

## Efficient Marginalization-Based MCMC Methods for Hierarchical Bayesian Inverse Problems\*

Arvind K. Saibaba<sup>†</sup>, Johnathan Bardsley<sup>‡</sup>, D. Andrew Brown<sup>§</sup>, and Alen Alexanderian<sup>†</sup>

**Abstract.** Hierarchical models in Bayesian inverse problems are characterized by an assumed prior probability distribution for the unknown state and measurement error precision, and hyper-priors for the prior parameters. Combining these probability models using Bayes' law often yields a posterior distribution that cannot be sampled from directly, even for a linear model with Gaussian measurement error and Gaussian prior, both of which we assume in this paper. In such cases, Gibbs sampling can be used to sample from the posterior [Bardsley, *SIAM J. Sci. Comput.*, 34 (2012), pp. A1316–A1332], but problems arise when the dimension of the state is large. This is because the Gaussian sample required for each iteration can be prohibitively expensive to compute, and because the statistical efficiency of the Markov chain degrades as the dimension of the state increases. The latter problem can be mitigated using marginalization-based techniques, such as those found in [Fox and Norton, *SIAM/ASA J. Uncertain. Quantif.*, 4 (2016), pp. 1191–1218; Joyce, Bardsley, and Luttman, *SIAM J. Sci. Comput.*, 40 (2018), pp. B766–B787; Rue and Held, *Monogr. Statist. Appl. Probab.* 104, Chapman & Hall/CRC, Boca Raton, FL, 2005], but these can be computationally prohibitive as well. In this paper, we combine the low-rank techniques of [Brown, Saibaba, and Valléian, *SIAM/ASA J. Uncertain. Quantif.*, 6 (2018), pp. 1076–1100] with the marginalization approach of [Rue and L. Held, *Monogr. Statist. Appl. Probab.* 104, Chapman & Hall/CRC, Boca Raton, FL, 2005]. We consider two variants of this approach: delayed acceptance and pseudomarginalization. We provide a detailed analysis of the acceptance rates and computational costs associated with our proposed algorithms and compare their performances on two numerical test cases—image deblurring and inverse heat equation.

**Key words.** Markov chain Monte Carlo, hierarchical Bayesian approach, inverse problems, one-block algorithm, low-rank approximations

**AMS subject classifications.** 65F22, 65C05, 65C60

**DOI.** 10.1137/18M1220625

**1. Introduction.** Inverse problems arise in a wide range of applications and typically involve estimating unknown parameters in a physical model from noisy, indirect measurements. In the applications that we are interested in, the unknown parameter vector results from the discretization of a continuous function defined on the computational domain and hence is

\*Received by the editors October 15, 2018; accepted for publication (in revised form) July 18, 2019; published electronically September 12, 2019.

<https://doi.org/10.1137/18M1220625>

**Funding:** This material was based upon work partially supported by the National Science Foundation under grant DMS-1638521 to the Statistical and Applied Mathematical Sciences Institute. The work of the first author was also partially supported by NSF DMS-1720398. The work of the third author was partially supported by NSF grants CMMI-1563435, EEC-1744497, and OIA-1826715.

<sup>†</sup>Department of Mathematics, North Carolina State University, Raleigh, NC 27695 (asaibab@ncsu.edu, alexanderian@ncsu.edu).

<sup>‡</sup>Department of Mathematical Sciences, University of Montana, Missoula, MT 59812 (bardsleyj@mso.umt.edu).

<sup>§</sup>School of Mathematical and Statistical Sciences, Clemson University, Clemson, SC 29634 (ab7@clemson.edu).

high-dimensional. Additionally, the inverse problem of recovering the unknown parameters from the data is *ill-posed*: a solution may not exist, may not be unique, or may be sensitive to the noise in the data.

To set the context of our work and some notation, we consider a discrete measurement error model of the form

$$(1) \quad \mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mu^{-1} \mathbf{I}_M),$$

where  $\mathbf{b} \in \mathbb{R}^M$  denotes the measurements,  $\mathbf{A} \in \mathbb{R}^{M \times N}$  is the discretized forward model,  $\mathbf{x} \in \mathbb{R}^N$  is the unknown estimand, and  $\mu > 0$  is the error precision. The inverse problem here seeks to recover the unknown  $\mathbf{x}$  from the measurement  $\mathbf{b}$ .

The statistical model (1) implies that the probability density function for  $\mathbf{b}$  is given by

$$(2) \quad \pi(\mathbf{b} \mid \mathbf{x}, \mu) \propto \mu^{M/2} \exp\left(-\frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2\right),$$

where “ $\propto$ ” denotes proportionality. To address the ill-posedness, we also assume a Gaussian prior probability density function (akin to choosing a quadratic regularization function) of the form

$$(3) \quad \pi(\mathbf{x} \mid \sigma) \propto \sigma^{N/2} \exp\left(-\frac{\sigma}{2} \mathbf{x}^\top \boldsymbol{\Gamma}_{\text{prior}}^{-1} \mathbf{x}\right).$$

For convenience, we take the prior to have zero mean, but a nonzero mean can also be easily incorporated into our framework. Then, if we define  $\boldsymbol{\theta} = (\mu, \sigma)$ , through *Bayes’ law* we obtain the *posterior density function* of  $\mathbf{x}$  conditioned on  $\mathbf{b}$  and  $\boldsymbol{\theta}$ :

$$(4) \quad \begin{aligned} \pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta}) &\propto \pi(\mathbf{b} \mid \mathbf{x}, \mu) \pi(\mathbf{x} \mid \sigma) \\ &\propto \mu^{M/2} \sigma^{N/2} \exp\left(-\frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 - \frac{\sigma}{2} \mathbf{x}^\top \boldsymbol{\Gamma}_{\text{prior}}^{-1} \mathbf{x}\right). \end{aligned}$$

The maximizer of  $\pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$  is known as the maximum a posteriori (MAP) estimator, which is also the minimizer of  $-\ln \pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$ . This, in turn, has the form of a Tikhonov regularization, thus establishing a connection between Bayesian and classical inverse problems.

The scaling terms in (4) involving  $\mu$  and  $\sigma$  arise from the normalizing constant of the Gaussian measurement error and prior models. In the hierarchical Bayesian approach we will also treat  $\boldsymbol{\theta} = (\mu, \sigma)$  as unknown. The a priori uncertainty about plausible values of  $\boldsymbol{\theta}$  is quantified in the prior distribution with density  $\pi(\boldsymbol{\theta})$ . For convenience, we assume that the precision parameters  $\mu$  and  $\sigma$  are independent and Gamma-distributed. Specifically, we assume

$$(5) \quad \pi(\boldsymbol{\theta}) = \pi(\mu) \pi(\sigma) \propto \mu^{\alpha_\mu - 1} \sigma^{\alpha_\sigma - 1} \exp(-\beta_\mu \mu - \beta_\sigma \sigma).$$

Other choices are also possible and are discussed in [24, 46, 8] and elsewhere. Applying Bayes’

law again yields a posterior density function of  $\mathbf{x}$  and  $\boldsymbol{\theta}$  conditioned on  $\mathbf{b}$ :

$$(6) \quad \pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) \propto \pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})\pi(\boldsymbol{\theta}).$$

The marginal distribution  $\pi(\boldsymbol{\theta} \mid \mathbf{b})$  can be derived by integrating over  $\mathbf{x}$ , i.e.,  $\pi(\boldsymbol{\theta} \mid \mathbf{b}) = \int_{\mathbb{R}^n} \pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) d\mathbf{x}$ . Explicit expressions for these distributions are provided in [section 2](#).

The focus of this paper is on the problem of sampling from the posterior distribution  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . In general, the full-posterior is not Gaussian and to explore this distribution the prevalent approach is to use Markov chain Monte Carlo (MCMC) methods [40, 30, 26, 50, 23]. For a more comprehensive review of MCMC methods, please refer to [6]. Several MCMC algorithms have been treated in the context of inverse problems in recent literature, which we briefly review. Specifically, in [3], conjugacy relationships are exploited to define a Gibbs sampler, in which samples from the conditional densities  $\pi(\boldsymbol{\theta} \mid \mathbf{b}, \mathbf{x})$  and  $\pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$  (which are Gamma- and Gaussian-distributed, respectively) are cyclically computed. The computational cost of this Gibbs sampler is prohibitive for  $N$  sufficiently large. This is due to the fact that as  $N \rightarrow \infty$ , the *integrated autocorrelation time* of the MCMC chain also tends to  $\infty$ , meaning that the number of Gibbs samples must increase with  $N$ ; see [1] for details. And second, computing samples from  $\pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$  requires the solution of an  $N \times N$  linear system, and hence the computational cost of the individual Gibbs samples also increases with  $N$ .

To address the first computational issue, i.e., that the correlation in the MCMC chain increases with  $N$ , an alternative to the Gibbs sampler is presented in [45], where a proposed state  $(\mathbf{x}_*, \boldsymbol{\theta}_*)$  is computed in two stages by first drawing  $\boldsymbol{\theta}_*$  from a proposal distribution and then drawing  $\mathbf{x}_* \sim \pi(\mathbf{x} \mid \boldsymbol{\theta}_*, \mathbf{b})$ . The proposal  $(\mathbf{x}_*, \boldsymbol{\theta}_*)$  is accepted or rejected jointly using a Metropolis–Hastings step to obtain an approximate draw from the posterior  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . This approach, called the *one-block* algorithm [45], does not have the same degeneracy issues as the Gibbs sampler as  $N \rightarrow \infty$ . However, it can be expensive to implement when evaluating  $\pi(\boldsymbol{\theta} \mid \mathbf{b})$  is computationally demanding and one still has to compute a sample from  $\pi(\mathbf{x} \mid \boldsymbol{\theta}_*, \mathbf{b})$  at every iteration.

To address the second computational issue, i.e., that the cost of computing samples from  $\pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$  increases with  $N$ , we implement the approach taken in [8], where a low-rank approximation of the so-called prior preconditioned data misfit part of the Hessian is used. This low-rank representation allows efficient sampling from the conditional distribution, reducing the overall computational cost. When the forward operator is defined using partial differential equations, computing the conditional covariance matrix once may require hundreds of thousands of PDE solves; in the context of an MCMC algorithm which requires repeated computation of the conditional covariance, this can be prohibitively expensive. Even when the conditional covariance can be formed, storing it requires  $\mathcal{O}(N^2)$  in memory and  $\mathcal{O}(N^3)$  in computational cost, which is infeasible when  $N$  is large (e.g.,  $\mathcal{O}(10^5)$ ). When the forward operator  $\mathbf{A}$  and the prior covariance matrix are diagonalized by the Fourier transform, algorithms such as the fast Fourier transform can be used effectively [5]. Other methods based on Krylov subspace solvers, e.g., conjugate gradient, have also been developed [27]. All of these approaches still suffer from the degeneracy issue as  $N \rightarrow \infty$ .

The contributions of this paper are as follows. First, to tackle the two drawbacks of the Gibbs sampler, as described above, we combine the approaches of the previous two paragraphs, which to our knowledge has not been done elsewhere. The use of a low-rank approximation

defines an approximate posterior density function  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ , whose samples are only approximate and thus must be embedded within a Metropolis–Hastings or importance sampling framework. In our first algorithm, which we call *approximate one-block* (AOB),  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is used as a proposal for Metropolis–Hastings, with the proposal samples computed using the one-block algorithm. We propose two other variants of the one-block algorithm that make use of  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . Specifically, we embed one-block applied to  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  within both the *delayed acceptance* [13] and *pseudomarginal* [2] frameworks to obtain samples from the full posterior  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . Thus the algorithms we propose result from combining the low-rank approximation approach of [8] with one of the existing MCMC methods mentioned above. To increase the novelty of our work, and also to provide the user with some intuition on how well our algorithms can be expected to perform in practice, we present theoretical results that provide insight into the acceptance rates and the performance of the algorithms. The main takeaway is that when the low-rank approximation is sufficiently close (in a sense that we make precise in section 4), the algorithms have similar behavior to the one-block algorithm.

Hierarchical Bayesian approaches have been applied to inverse problems in various other works, going back to [35] and more recently in [11, 12]. However, in those works the posterior density function  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is maximized, yielding the MAP estimator, whereas in this paper we want to perform uncertainty quantification (UQ), and so need to compute samples from the posterior. Sample-based methods for inverse problems first appear in the works [41, 34]. In recent years, MCMC methods for Bayesian inverse problems has become an active field, with some recent advances including gradient and Hessian-based MCMC methods [39, 43], likelihood-informed MCMC methods [17], and transport map accelerated MCMC methods [42]. In the Bayesian statistics literature, MCMC methods for hierarchical models of the type considered here are standard; see, e.g., [25]. Moreover, the Gibbs sampler of [3] for inverse problems is used in the context of spatio-temporal models in [31]. Some properties of this Gibbs sampler are derived in [1], and various extensions are presented in [8, 22, 33], which have improved convergence properties and/or improved computational efficiency. The algorithms presented in this paper fit within this last group of MCMC methods.

The paper is organized as follows. In section 2, we present the hierarchical Gibbs sampler and discuss the infinite-dimensional limit, presenting the result of [1] showing the degeneracy of hierarchical Gibbs as  $N \rightarrow \infty$ . We then present the one-block algorithm of [45], which does not have the same degeneracy issues. In section 3, we present the new algorithms making only limited assumptions regarding the approximate full posterior distribution. In section 4, we provide a theoretical analysis of our proposed algorithms, making explicit use of low-rank structure. The numerical experiments in section 5 include a model one-dimensional deblurring problem which allows us to compare and contrast the various algorithms and a PDE-based example of inverse heat equation that demonstrates the computational benefits of our approaches. We summarize our work and discuss future research in section 6.

**2. Review of MCMC algorithms.** In this section, we present two MCMC algorithms for background. The first method is known as hierarchical Gibbs and is standard. Its convergence characteristics serve as motivation for the second algorithm, which is known as one-block.

**2.1. The hierarchical Gibbs sampler.** Under our assumed model, we provide explicit expressions for the posterior and the marginal distribution. To obtain the expression for the

posterior distribution, combine (4) with (5) via (6) to obtain

$$\begin{aligned}\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) &\propto \pi(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})\pi(\boldsymbol{\theta}) \\ &\propto \mu^{M/2}\sigma^{N/2}\pi(\boldsymbol{\theta}) \exp\left(-\frac{\mu}{2}\|\mathbf{Ax} - \mathbf{b}\|^2 - \frac{\sigma}{2}\mathbf{x}^\top \boldsymbol{\Gamma}_{\text{prior}}^{-1}\mathbf{x}\right) \\ &\propto \mu^{M/2}\sigma^{N/2}\pi(\boldsymbol{\theta}) \exp\left(-\frac{\mu}{2}\mathbf{b}^\top \mathbf{b} + \mu\mathbf{b}^\top \mathbf{Ax} - \frac{1}{2}\mathbf{x}^\top \boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta})\mathbf{x}\right) \\ &= \mu^{M/2}\sigma^{N/2}\pi(\boldsymbol{\theta}) \exp\left(-\frac{\mu}{2}\mathbf{b}^\top \mathbf{b} + \frac{\mu^2}{2}\mathbf{b}^\top \mathbf{A}\boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta})\mathbf{A}^\top \mathbf{b}\right) \\ &\quad \times \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}_{\text{cond}}(\boldsymbol{\theta}))^\top \boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta})(\mathbf{x} - \mathbf{x}_{\text{cond}}(\boldsymbol{\theta}))\right),\end{aligned}$$

where  $\mathbf{x}_{\text{cond}}(\boldsymbol{\theta}) = \mu\boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta})\mathbf{A}^\top \mathbf{b}$  and  $\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) = \mu\mathbf{A}^\top \mathbf{A} + \sigma\boldsymbol{\Gamma}_{\text{prior}}^{-1}$ . It follows that the marginal distribution,  $\pi(\boldsymbol{\theta} \mid \mathbf{b}) = \int_{\mathbb{R}^n} \pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})d\mathbf{x}$ , is given by

$$(7) \quad \pi(\boldsymbol{\theta} \mid \mathbf{b}) \propto \frac{\mu^{M/2}\sigma^{N/2}\pi(\boldsymbol{\theta})}{\sqrt{\det(\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}))}} \exp\left(-\frac{\mu}{2}\mathbf{b}^\top \mathbf{b} + \frac{\mu^2}{2}\mathbf{b}^\top \mathbf{A}\boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta})\mathbf{A}^\top \mathbf{b}\right).$$

Observe that the joint posterior density satisfies  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) = \pi(\boldsymbol{\theta} \mid \mathbf{b})\pi(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{b})$ , where  $\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{b} \sim \mathcal{N}(\mathbf{x}_{\text{cond}}(\boldsymbol{\theta}), \boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta}))$ .

We begin with the hierarchical Gibbs sampler of [3]. Our choice of prior (3) for  $\mathbf{x}$ , and the hyper-prior (5) for  $\boldsymbol{\theta} = (\mu, \sigma)$ , respectively, were made with conjugacy relationships in mind [25], i.e., so that the “full conditional” densities have the same form as the corresponding priors:

$$(8) \quad \pi(\mu \mid \mathbf{x}, \sigma, \mathbf{b}) \propto \mu^{M/2+\alpha_\mu-1} \exp\left(\left[-\frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|_2^2 - \beta_\mu\right]\mu\right),$$

$$(9) \quad \pi(\sigma \mid \mathbf{x}, \mu, \mathbf{b}) \propto \sigma^{N/2+\alpha_\sigma-1} \exp\left(\left[-\frac{1}{2}\mathbf{x}^\top \boldsymbol{\Gamma}_{\text{prior}}^{-1}\mathbf{x} - \beta_\sigma\right]\sigma\right),$$

$$(10) \quad \pi(\mathbf{x} \mid \mu, \sigma, \mathbf{b}) \propto \exp\left(-\frac{\mu}{2}\|\mathbf{Ax} - \mathbf{b}\|^2 - \frac{\sigma}{2}\mathbf{x}^\top \boldsymbol{\Gamma}_{\text{prior}}^{-1}\mathbf{x}\right).$$

Note that (8) and (9) are Gamma densities, while (10) is the density of a Gaussian distribution. Algorithm 1 follows immediately from (8)–(10) and is precisely the hierarchical Gibbs sampling algorithm of [3].

The values of  $M$  and  $N$  are determined by the number of measurements and the size of the numerical mesh, respectively, making the problems discrete. Since  $M$  is the dimension of our measurement vector, we assume that it is a fixed value. However, we are free to choose  $N$  as we please, and the behavior of our approaches as  $N \rightarrow \infty$  is an important question. In what follows, we briefly discuss the infinite-dimensional limit, pointing the interested reader to the extensive treatments found in [49, 20] for more details.

Consider the linear inverse problem, which typically arises from the discretization of a Fredholm integral equation of the first kind, for example,

$$(11) \quad b(s) = \int_{\Omega} a(s; t)x(t)dt, \quad s \in \Omega,$$

**Algorithm 1** Hierarchical Gibbs Sampler.

**Input:** Set  $\mathbf{x}_{(0)} = \mathbf{x}_{\text{cond}}(\boldsymbol{\theta}_{(0)})$ , and define  $K$  and burn-in period  $K_b$ .

**Output:** Approximate samples from the posterior distribution  $\{\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}\}_{t=K_b+1}^K$ .

---

```

1 for  $t = 1$  to  $K$  do
2   |   Compute  $\mu_{(t)} \sim \Gamma(M/2 + \alpha_\mu, \frac{1}{2}\|\mathbf{A}\mathbf{x}_{(t-1)} - \mathbf{b}\|^2 + \beta_\mu)$ .
3   |   Compute  $\sigma_{(t)} \sim \Gamma(N/2 + \alpha_\sigma, \frac{1}{2}(\mathbf{x}_{(t-1)})^T \boldsymbol{\Gamma}_{\text{prior}} \mathbf{x}_{(t-1)} + \beta_\sigma)$ .
4   |   Compute  $\mathbf{x}_{(t)} \sim \mathcal{N}(\mathbf{x}_{\text{cond}}(\boldsymbol{\theta}_{(t)}), \boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta}_{(t)}))$ , where  $\boldsymbol{\theta}_{(t)} = (\mu_{(t)}, \sigma_{(t)})$ .
5 end

```

---

where  $b$  is the model output function,  $\Omega$  is the computational domain,  $a$  is the integral kernel or point spread function, and  $x$  is the unknown which we seek to estimate. We define  $\mathcal{A}_M x$  to be the forward operator discretized only in its range, so that  $\mathcal{A}_M : X \rightarrow \mathbb{R}^M$ , where  $X = C(\bar{\Omega})$ , and  $\Omega$  is the spatial domain. For example, in one-dimensional deconvolution, (with  $\Omega = (0, 1)$ ) one can have

$$[\mathcal{A}_M x]_i = \int_0^1 a(s_i - s') x(s') ds', \quad i = 1, \dots, M.$$

Discretizing this integral in the  $s'$  variable, e.g., using a uniform mesh on  $[0, 1]$  with  $N$  grid elements and midpoint quadrature, then yields  $\mathbf{A}\mathbf{x}$ . Then we have that

$$(12) \quad \lim_{N \rightarrow \infty} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 = \|\mathcal{A}_M x - \mathbf{b}\|_2^2.$$

Note that here  $\mathbf{x}$  denotes the discretized version of  $x \in X$ . For the prior (3), it is typical to choose  $\boldsymbol{\Gamma}_{\text{prior}}$  to be the numerical discretization of the inverse of a differential operator. That is, letting  $\Gamma_{\text{prior}}$  denote the infinite-dimensional prior covariance operator, we define  $\boldsymbol{\Gamma}_{\text{prior}} = \mathcal{L}^{-1}$ , where  $\mathcal{L}$  is a differential operator. A basic requirement on  $\boldsymbol{\Gamma}_{\text{prior}}$  is that it has to be trace-class on  $L^2(\Omega)$ . That is, for any orthonormal basis  $\{\phi_i\}_{i=1}^\infty$  of  $L^2(\Omega)$ ,  $\sum_{i=1}^\infty \langle \phi_i, \mathcal{L}^{-1} \phi_i \rangle < \infty$ ; see [49]. Moreover, if  $\langle x, y \rangle$  is the standard  $L^2(\Omega)$  inner product, we can use midpoint quadrature to obtain  $\langle x, y \rangle = \lim_{N \rightarrow \infty} \langle \mathbf{x}, \mathbf{y} \rangle_N$ , where the boldface letters indicate discretized versions of the variables and  $\langle \mathbf{x}, \mathbf{y} \rangle_N := \frac{1}{N} \sum_{i=1}^N x_i y_i$ . Using this notation,

$$(13) \quad \lim_{N \rightarrow \infty} \langle \mathbf{x}, \boldsymbol{\Gamma}_{\text{prior}}^{-1} \mathbf{x} \rangle_N = \langle x, \mathcal{L} x \rangle.$$

Combining (12) and (13) yields

$$(14) \quad \lim_{N \rightarrow \infty} \left\{ \frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\sigma}{2} \langle \mathbf{x}, \boldsymbol{\Gamma}_{\text{prior}}^{-1} \mathbf{x} \rangle_N \right\} = \frac{\mu}{2} \|\mathcal{A}_M x - \mathbf{b}\|_2^2 + \frac{\sigma}{2} \langle x, \mathcal{L} x \rangle.$$

There are two issues that arise when we consider the limit as  $N \rightarrow \infty$ : the first is mathematical, and the second is computational. A question that immediately arises is whether or not one can define an infinite-dimensional limit of the posterior density function. We cannot define a probability density function on a function space. The reason for this is that in finite dimensions, the posterior density is none but the Radon–Nikodym derivative of



posterior probability law of the inference parameter with respect to Lebesgue measure, but one cannot define a Lebesgue measure on an infinite-dimensional function space. However, as shown in [49], given that the prior covariance operator is trace class and  $\mathcal{A}_M : X \rightarrow \mathbb{R}^M$  is a continuous linear transformation, the posterior law of  $x$ , which we denote by  $\nu^{\mathbf{b}, \boldsymbol{\theta}}$ , is a Gaussian measure on  $L^2(\Omega)$ ,  $\nu^{\mathbf{b}, \boldsymbol{\theta}} = \mathcal{N}(x_{\text{cond}}(\boldsymbol{\theta}), \Gamma_{\text{cond}}(\boldsymbol{\theta}))$ , with

$$\Gamma_{\text{cond}}^{-1} = \mu \mathcal{A}_M^* \mathcal{A}_M + \sigma \Gamma_{\text{prior}}^{-1}, \quad x_{\text{cond}} = \mu \Gamma_{\text{cond}} \mathcal{A}_M^* \mathbf{b}.$$

As for the construction of the prior covariance operator, as mentioned before, a common approach is to define them as inverses of differential operators. For example, we can define

$$\Gamma_{\text{prior}} = \mathcal{L}^{-s}, \quad \mathcal{L}u = -\kappa \Delta u + \alpha u, \quad \kappa > 0, \alpha \geq 0,$$

with suitable boundary conditions, which is related to the Whittle–Matérn prior [4]. In this paper, we choose  $\alpha = 0$ ,  $\kappa = \sigma^{-1/s}$ , and homogeneous Dirichlet boundary conditions for the Laplacian. Therefore,  $\Gamma_{\text{prior}} = \sigma(-\Delta)^{-s}$ , leaving only  $\sigma$  as the hyper-parameter for the prior. Moreover, to ensure that the covariance operator so defined is trace-class, we require  $s > d/2$ , where  $d$  is the spatial dimension of the problem. For the one-dimensional example in (11),  $s = 1$  would suffice. For problems with  $d = 2$  or  $d = 3$ , a convenient option is  $s = 2$ . Note this assumption on  $s$  also ensures that the prior draws are almost surely continuous. For further details on the definition of Gaussian measures on infinite-dimensional Hilbert spaces, see [18, 19].


A second question that arises is whether or not the performance of the hierarchical Gibbs sampler is dependent upon  $N$ . For inverse problems in which the infinite-dimensional limit is well-defined, MCMC methods whose performance is independent of the discretization ( $N$  in this case) are desirable. In line 3 of Algorithm 1, we see that  $N$  appears in the Gamma conditional density  $\pi(\sigma \mid \mathbf{x}, \mu, \mathbf{b})$ , and thus it should not be surprising to find the  $\sigma$ -chain is dependent on  $N$ . The exact nature of this dependence is the subject of [1, Theorem 3.4], where under reasonable assumptions it is shown that the expected step in the  $\sigma$ -chain scales like  $2/N$ . Specifically, for any  $\sigma > 0$ ,

$$\frac{N}{2} \mathbb{E} [\sigma_{(t+1)} - \sigma_{(t)} \mid \sigma_{(t)} = \sigma] = (\alpha_\sigma + 1)\sigma - f_N(\sigma; \mathbf{b})\sigma^2 + \mathcal{O}(N^{-1/2}),$$

where  $\mathbb{E}$  denotes expectation and  $f_N(\sigma; \mathbf{b})$  is bounded uniformly in  $N$ . Moreover, the variance of the step also scales like  $2/N$ ; for any  $\sigma > 0$ ,

$$\frac{N}{2} \text{Var} [\sigma_{(t+1)} - \sigma_{(t)} \mid \sigma_{(t)} = \sigma] = 2\sigma^2 + \mathcal{O}(N^{-1/2}).$$

A consequence of these results, as is noted in [1], is that the expected squared jumping distance of the Markov chain for  $\sigma$  is  $\mathcal{O}(1/N)$ . Moreover, it is noted that the lag-1 autocorrelation of the  $\sigma$ -chain behaves like  $1 - c/N$  for some constant, but  $\text{Var}(\sigma_{(t)}) = \mathcal{O}(1)$ . Hence, the Monte Carlo error associated with  $K - K_b$  draws in stationarity is  $\mathcal{O}(\sqrt{N/(K - K_b)})$ . Thus, the  $\sigma$ -chain becomes increasingly correlated as  $N \rightarrow \infty$ . This phenomenon is illustrated in Figure 1, which displays the empirical autocorrelation functions for the  $\mu$ - and  $\sigma$ -chains generated by Algorithm 1 for a one-dimensional image deblurring test problem. Note that as  $N$  increases by a power of 2, so does the integrated autocorrelation time (IACT).



**Figure 1.** Plots of the empirical autocorrelation functions for the  $\sigma$  chains in the one-dimensional deblurring example: (left) block-Gibbs and (right) AOB discussed in [subsection 3.1](#). See [subsection 5.1](#) for details regarding the experimental setup.

**2.2. One-block MCMC.** An alternative algorithm for computing a sample  $(\mathbf{x}', \boldsymbol{\theta}') \sim \pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})$  is to first compute a sample from the marginal density  $\boldsymbol{\theta}' \sim \pi(\boldsymbol{\theta}' | \mathbf{b})$ , defined in (7), and then compute a sample from the conditional density  $\mathbf{x}' \sim \pi(\mathbf{x}' | \mathbf{b}, \boldsymbol{\theta}')$ . In principle, we can use this procedure as an alternative to [Algorithm 1](#). In practice, it is often not possible to sample directly from  $\pi(\boldsymbol{\theta} | \mathbf{b})$ . In the *one-block algorithm* of [45], a Markov chain is generated in which if  $(\mathbf{x}, \boldsymbol{\theta})$  is the current element of the chain, a state  $\boldsymbol{\theta}'$  is proposed from some proposal distribution with density  $r(\boldsymbol{\theta}' | \boldsymbol{\theta})$ , then  $\mathbf{x}' \sim \pi(\mathbf{x}' | \mathbf{b}, \boldsymbol{\theta}')$  is drawn, and the pair  $(\mathbf{x}', \boldsymbol{\theta}')$  is accepted with probability

$$\begin{aligned} \rho_1(\boldsymbol{\theta}'; \boldsymbol{\theta}) &= \min \left\{ 1, \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})} \frac{r(\boldsymbol{\theta} | \boldsymbol{\theta}') \pi(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})}{r(\boldsymbol{\theta}' | \boldsymbol{\theta}) \pi(\mathbf{x}' | \mathbf{b}, \boldsymbol{\theta}')} \right\} \\ (15) \quad &= \min \left\{ 1, \frac{\pi(\boldsymbol{\theta}' | \mathbf{b}) r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta} | \mathbf{b}) r(\boldsymbol{\theta}' | \boldsymbol{\theta})} \right\}. \end{aligned}$$

The resulting MCMC method is given by [Algorithm 2](#).

---

**Algorithm 2** One-Block MCMC.

---

**Input:** Set  $\mathbf{x}_{(0)} = \mathbf{x}_{\text{cond}}(\boldsymbol{\theta}_{(0)})$ , and define  $K$  and burn-in period  $K_b$ .

**Output:** Approximate samples from the posterior distribution  $\{\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}\}_{t=K_b+1}^K$ .

---

```

1 for  $t = 1$  to  $K$  do
2   Compute  $\boldsymbol{\theta}_* \sim r(\boldsymbol{\theta}_* | \boldsymbol{\theta}_{(t-1)})$ .
3   Compute  $\mathbf{x}_{(t)} \sim \pi(\mathbf{x}_{(t)} | \mathbf{b}, \boldsymbol{\theta}_*)$ .
4   Set  $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}_*, \boldsymbol{\theta}_*)$  with probability  $\rho_1(\boldsymbol{\theta}_*; \boldsymbol{\theta}_{(t-1)})$ , defined in (15), else set
       $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$ .
5 end
```

---



Note that by combining the random walk  $\theta' \sim r(\theta' | \theta)$  with the conditional sample  $x' \sim \pi(x' | b, \theta')$ ,  $x$  is marginalized (i.e., integrated out) from the posterior, as is seen by the acceptance ratio (15) which is independent of  $x$ . Note that if  $r(\cdot | \cdot)$  is used as a Metropolis–Hastings proposal for sampling from the marginal density  $\pi(\theta | b)$  (which defines the MCMC method proposed in [22]), the acceptance ratio is also given by (15). Thus, the  $\theta$ -chain generated by Algorithm 2 converges in distribution to  $\pi(\theta | b)$  and its behavior is independent of  $x$ , as desired.

As an alternative to MCMC, we remark that there exist special cases of the hierarchical Bayesian model in which the hyperparameters  $\theta$  can be analytically integrated out to obtain a closed-form marginal posterior density  $x | b$  (e.g., multivariate  $t$ ). Despite the fact that MCMC is no longer necessary in such scenarios, there is still the need to work with a scale matrix that may require many forward solutions to compute, rendering such an approach impractical. Flath et al. [21] study this issue in the Gaussian case in which the precision parameters are held fixed (as opposed to integrating them out). The same problems arise with any other distribution of  $x | b$ , but with additional difficulties since other forms are generally no longer in the Gaussian class of distributions. Further, a researcher may wish to use alternative hyperpriors other than those that lead to closed-form distributions for the sake of obtaining a better reconstruction. The approach we consider here is flexible with respect to such considerations because the conditional distribution of  $x | \theta, b$  is still Gaussian, regardless of the prior on  $\theta$ . Exploration of alternative priors for  $\theta$  is beyond the scope of the present work.

**3. Algorithms.** For some applications, evaluating the marginal density  $\pi(\theta | b)$  is not computationally tractable, nor is computing exact samples  $x' \sim \pi(x' | b, \theta')$ . In this case, implementation of Algorithm 2 is infeasible. However, suppose that we have an approximate posterior density function  $\hat{\pi}(x, \theta | b)$  for which the marginal density  $\hat{\pi}(\theta | b)$  can be evaluated, and exact samples  $x' \sim \hat{\pi}(x' | b, \theta')$  can be drawn efficiently.

In this section, we propose three MCMC algorithms, each of which uses  $\hat{\pi}(x, \theta | b)$  to approximate Algorithm 2. We note that Rue and Held [45] also consider approximate marginalization of  $x$  in the one-block algorithm. However, the difference between our approach and that of [45] is the nature of the approximate posterior distribution. Specifically, motivated by non-Gaussian conditional distributions that arise in, e.g., Poisson counts for disease mapping, Rue and Held [45] use an approximation based on a Taylor expansion of the log-likelihood about the mode of the full conditional distribution, resulting in a Gaussian proposal in the Metropolis–Hastings algorithm. By contrast, we are concerned with cases in which the full conditional is still Gaussian distributed, but drawing realizations from the distribution and evaluating the density are computationally prohibitive because of the dimensionality of the unknown parameters. We discuss in subsection 4.4 our construction of the approximation  $\hat{\pi}(x, \theta | b)$  and its effect on the proposed algorithms.

**3.1. Approximate one-block MCMC.** Suppose we generate a proposal  $(x', \theta')$  by first drawing  $\theta'$  from a proposal distribution with density  $r(\theta' | \theta)$  followed by drawing  $x' \sim \hat{\pi}(x' | \theta', b)$ . Using this in a Metropolis–Hastings sampler with target distribution  $\pi(x, \theta | b)$  is simply the one-block algorithm with  $\hat{\pi}(x | \theta, b)$  in place of  $\pi(x | \theta, b)$ . The proposal density is then given by

$$q(x', \theta' | x, \theta) = \hat{\pi}(x' | \theta', b) r(\theta' | \theta),$$

**Algorithm 3** Approximate One-Block MCMC.

**Input:** Initialize  $\boldsymbol{\theta}_{(0)}$ . Define the length of the chain  $C$  and burn-in period  $C_b$ .

**Output:** Approximate sample from the posterior distribution  $\{\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}\}_{t=C_b+1}^K$ .

---

```

1 Compute  $\mathbf{x}_{(0)} \sim \hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta}_{(0)})$ .
2 for  $t = 1$  to  $K$  do
3   Draw  $\boldsymbol{\theta}'$  from distribution with density  $r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)})$ .
4   Draw  $\mathbf{x}'$  from distribution with density  $\hat{\pi}(\mathbf{x}' \mid \boldsymbol{\theta}', \mathbf{b})$ .
5   Set  $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}', \boldsymbol{\theta}')$  with probability  $\rho_2(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$  defined in (16); else
      set  $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$ .
6 end

```

---

and hence, the acceptance probability is

$$(16) \quad \rho_2(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{b}) \hat{\pi}(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{b}) r(\boldsymbol{\theta} \mid \boldsymbol{\theta}')}{\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) \hat{\pi}(\mathbf{x}' \mid \boldsymbol{\theta}', \mathbf{b}) r(\boldsymbol{\theta}' \mid \boldsymbol{\theta})} \right\}.$$

The full procedure is summarized in Algorithm 3.

There are two observations to make about Algorithm 3, each of which motivates the MCMC methods that follow. First, computing the ratio (16) requires evaluating the true posterior density  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . In cases in which this is computationally burdensome, the MCMC method presented next seeks to avoid extraneous evaluations of  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  using a technique known as *delayed acceptance* [13]. Second, Algorithm 3 only *approximately* integrates  $\mathbf{x}$  out of  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ , since  $\pi(\boldsymbol{\theta} \mid \mathbf{b}) \approx \pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) / \hat{\pi}(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{b})$ , but this is not an exact marginalization and thus some dependence between the  $\mathbf{x}$  and  $\boldsymbol{\theta}$  chains may remain, slowing convergence of the algorithm. However, as we discuss below, Algorithm 3 is a special case of the so-called *pseudomarginal MCMC algorithm* [2]. By generalizing the AOB algorithm, we can improve the approximation to  $\pi(\boldsymbol{\theta} \mid \mathbf{b})$  and thus improve the mixing of the Markov chains.

**3.2. Approximate one-block MCMC with delayed acceptance.** It will often be the case that evaluating  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is computationally prohibitive and/or significantly more expensive than evaluating  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ , in which case one would like to generate samples from  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  while minimizing the number of times its density is evaluated. This motivates the use of the *delayed acceptance* framework of [13] to improve the computational efficiency of Algorithm 3. In this approach, one step of Algorithm 2 is used with  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  taken as the target distribution. Only if  $(\mathbf{x}', \boldsymbol{\theta}')$  is accepted as a sample from  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is it proposed as a sample from  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ . The idea is to only evaluate  $\pi(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{b})$  for proposed states that have a high probability of being accepted as draws from the true distribution. This prevents us from wasting computational effort on rejected proposals. The approximate distribution in the first stage is essentially a computationally cheap “filter” that prevents this from occurring. The resulting *approximate one-block MCMC with delayed acceptance (ABDA)* procedure is given in Algorithm 4. Note that the delayed acceptance algorithm proposed here is related to the surrogate transition method [38, section 9.4.3] and is a special case of [13] with a state-independent approximation.

**Algorithm 4** Approximate One-Block MCMC with Delayed Acceptance.**Input:** Initialize  $\boldsymbol{\theta}_{(0)}$ . Define the length of the chain  $C$  and burn-in period  $C_b$ .**Output:** Approximate sample from the posterior distribution  $\{\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}\}_{t=C_b+1}^C$ .

```

1 Compute  $\mathbf{x}_{(0)} \sim \hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta}_{(0)})$ .
2 for  $t = 1$  to  $K$  do
    // Stage 1: apply the one-block algorithm to  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ .
3   Draw  $\boldsymbol{\theta}_*$  from the distribution with density  $r(\boldsymbol{\theta}_* \mid \boldsymbol{\theta}_{(t-1)})$ .
4   Draw  $\mathbf{x}_*$  from the distribution with density  $\hat{\pi}(\mathbf{x}_* \mid \boldsymbol{\theta}_*, \mathbf{b})$ .
5   Compute  $\hat{\rho}_1(\boldsymbol{\theta}_*; \boldsymbol{\theta}_{(t-1)})$  defined in (18).
6   Set  $(\mathbf{x}', \boldsymbol{\theta}') \leftarrow (\mathbf{x}_*, \boldsymbol{\theta}_*)$  and promote to Stage 2 with probability  $\hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)})$  else set
       $(\mathbf{x}', \boldsymbol{\theta}') \leftarrow (\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$  and promote.
    // Stage 2: accept/reject  $(\mathbf{x}', \boldsymbol{\theta}')$  as a proposal for  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ .
7   Define
      (17)  $\rho_3(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{b}) q(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{x}', \boldsymbol{\theta}')}{\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) q(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{x}, \boldsymbol{\theta})} \right\},$ 
      where  $q(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$  is defined in (19). Set  $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}', \boldsymbol{\theta}')$  with proba-
      bility  $\rho_3(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$ , else set  $(\mathbf{x}_{(t)}, \boldsymbol{\theta}_{(t)}) = (\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$ .
8 end

```

Let  $\hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)})$  be given by (15), but with  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  replaced by  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ , i.e.,

$$\begin{aligned}
 \hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)}) &= \min \left\{ 1, \frac{\hat{\pi}(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{b})}{\hat{\pi}(\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)} \mid \mathbf{b})} \frac{r(\boldsymbol{\theta}_{(t-1)} \mid \boldsymbol{\theta}') \hat{\pi}(\mathbf{x}_{(t-1)} \mid \mathbf{b}, \boldsymbol{\theta}_{(t-1)})}{r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)}) \hat{\pi}(\mathbf{x}' \mid \mathbf{b}, \boldsymbol{\theta}')} \right\} \\
 (18) \quad &= \min \left\{ 1, \frac{\hat{\pi}(\boldsymbol{\theta}' \mid \mathbf{b}) r(\boldsymbol{\theta}_{(t-1)} \mid \boldsymbol{\theta}')}{\hat{\pi}(\boldsymbol{\theta}_{(t-1)} \mid \mathbf{b}) r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)})} \right\}.
 \end{aligned}$$

Then, Stage 2 of Algorithm 4 is a Metropolis–Hastings algorithm with target distribution  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  and proposal density given by

$$\begin{aligned}
 (19) \quad q(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)}) &= \hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)}) \hat{\pi}(\mathbf{x}' \mid \boldsymbol{\theta}', \mathbf{b}) r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)}) \\
 &\quad + \delta_{(\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})}(\mathbf{x}', \boldsymbol{\theta}') (1 - y(\boldsymbol{\theta}_{(t-1)})),
 \end{aligned}$$

where  $y(\boldsymbol{\theta}_{(t-1)}) = \int \int \hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)}) \hat{\pi}(\mathbf{x}' \mid \boldsymbol{\theta}', \mathbf{b}) r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)}) d\mathbf{x}' d\boldsymbol{\theta}'$ . Note that there is never a need to evaluate  $y(\boldsymbol{\theta}_{(t-1)})$ , since when the promoted state is  $(\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$ ,

$$\rho_3(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)}) = 1,$$

and the chain remains at the same point. Conversely, when the composite sample  $(\mathbf{x}', \boldsymbol{\theta}')$  is promoted,  $(\mathbf{x}', \boldsymbol{\theta}') \neq (\mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)})$  and

$$q(\mathbf{x}', \boldsymbol{\theta}' \mid \mathbf{x}_{(t-1)}, \boldsymbol{\theta}_{(t-1)}) = \hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}_{(t-1)}) \hat{\pi}(\mathbf{x}' \mid \boldsymbol{\theta}', \mathbf{b}) r(\boldsymbol{\theta}' \mid \boldsymbol{\theta}_{(t-1)}).$$

We discuss in [section 4](#) the conditions under which the acceptance rate in the second stage is high, thus preventing wasted computational effort on rejected proposals.

The validity of the ABDA algorithm can be gleaned by recognizing that it is a special case of the delayed acceptance algorithm in [\[13\]](#), with first stage proposal given by  $\hat{\pi}(\mathbf{x} \mid \boldsymbol{\theta}', \mathbf{b})r(\boldsymbol{\theta}' \mid \boldsymbol{\theta})$ , the approximating density given by  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ , and target density  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{y})$ . This latter observation allows us to apply Theorem 1 of [\[13\]](#) to establish the irreducibility and aperiodicity of the ABDA Markov chain. Standard ergodic theory (e.g., [\[44\]](#)) then provides that  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is, in fact, the limiting distribution.

**3.3. Pseudomarginal MCMC.** As noted at the end of [Algorithm 2](#), the  $\boldsymbol{\theta}$ -chain generated by [Algorithm 2](#) is independent of the  $\mathbf{x}$ -chain. Hence, it does not suffer from the same degeneracy issues as the  $\boldsymbol{\theta}$ -chain generated by [Algorithm 1](#). When it is not computationally feasible to implement [Algorithm 2](#), and an approximate posterior  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is used as in [Algorithms 3](#) and [4](#), the marginalization is only approximate so that there still remains some dependence between the  $\boldsymbol{\theta}$  and  $\mathbf{x}$  chains. This dependence can be mitigated as the approximation to the marginal distribution of  $\boldsymbol{\theta}$  improves. In particular, we can use importance sampling, with importance density  $\hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$ , to approximate the integration over  $\mathbf{x}$ . Specifically, we have that

$$\begin{aligned}
 \pi(\boldsymbol{\theta}' \mid \mathbf{b}) &= \int_{\mathbb{R}^N} \pi(\mathbf{x}, \boldsymbol{\theta}' \mid \mathbf{b}) d\mathbf{x} \\
 &= \int_{\mathbb{R}^N} \frac{\pi(\mathbf{x}, \boldsymbol{\theta}' \mid \mathbf{b})}{\hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta}')} \hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta}') d\mathbf{x} \\
 (20) \quad &\approx \frac{1}{K} \sum_{j=1}^K \frac{\pi(\mathbf{x}'_j, \boldsymbol{\theta}' \mid \mathbf{b})}{\hat{\pi}(\mathbf{x}'_j \mid \mathbf{b}, \boldsymbol{\theta}')} \stackrel{\text{def}}{=} \pi^K(\boldsymbol{\theta}' \mid \mathbf{b}),
 \end{aligned}$$

where  $\mathbf{x}'_j \sim \hat{\pi}(\mathbf{x}' \mid \mathbf{b}, \boldsymbol{\theta}')$ . The idea behind pseudomarginal MCMC [\[2\]](#), in our setting, is to generalize [Algorithm 3](#) by using  $\pi^K(\boldsymbol{\theta} \mid \mathbf{b})$  as an approximation to  $\pi(\boldsymbol{\theta} \mid \mathbf{b})$ . The resulting algorithm is given in [Algorithm 5](#).

When  $K = 1$ , [Algorithm 5](#) simply reduces to [Algorithm 3](#). Conversely, as  $K \rightarrow \infty$ ,  $\pi^K(\boldsymbol{\theta}' \mid \mathbf{b}) \rightarrow \pi(\boldsymbol{\theta}' \mid \mathbf{b})$ , and hence the Markov chains produced by [Algorithms 2](#) and [5](#) become indistinguishable. Consequently, as  $K$  increases, the dependence between the  $\boldsymbol{\theta}$ -chain and the  $\mathbf{x}$ -chain dissipates as desired. It is shown in [\[2\]](#) that the value of  $\pi^K(\boldsymbol{\theta}_{(t-1)} \mid \mathbf{b})$  computed in step  $t - 1$  can be reused in step  $t$  so that a new set of importance samples does not need to be computed in [\(20\)](#). Furthermore, this paper showed that the corresponding Markov chain will converge in distribution to the target density, in our case  $\pi(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$ .

**4. Analysis.** [Algorithms 3](#) to [5](#) are all approximations of the one-block algorithm, [Algorithm 2](#), and they are meant to be used in cases in which implementing one-block is computationally expensive. We assume  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b})$  is given by the first equation array in [subsection 2.1](#) (array is not numbered), but with the conditional covariance  $\boldsymbol{\Gamma}_{\text{cond}}(\boldsymbol{\theta})$  replaced by an approximation  $\hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta})$ .

**Algorithm 5** Pseudomarginal MCMC.

**Input:** Initialize  $\boldsymbol{\theta}_{(0)}$ . Define the length of the chain  $C$  and burn-in period  $C_b$ .

**Output:** Approximate samples from the posterior distribution  $\{\{\mathbf{x}_{(t,i)}\}_{i=1}^K, \boldsymbol{\theta}_{(t)}\}_{t=C_b+1}^C$ .

```

1 Compute  $\pi^K(\boldsymbol{\theta}_{(0)} | \mathbf{b})$  with  $\mathbf{x}'_j \sim \hat{\pi}(\mathbf{x}' | \mathbf{b}, \boldsymbol{\theta}_{(0)})$  as in (20). Define  $\{\mathbf{x}_{(0,j)}\}_{j=1}^K = \{\mathbf{x}'_j\}_{j=1}^K$ .
2 for  $t = 1$  to  $C$  do
3   Compute  $\boldsymbol{\theta}' \sim r(\cdot | \boldsymbol{\theta}_{(t-1)})$ .
4   Compute  $\pi^K(\boldsymbol{\theta}' | \mathbf{b})$  using  $\mathbf{x}'_j \sim \hat{\pi}(\mathbf{x}' | \mathbf{b}, \boldsymbol{\theta}')$  as in (20).
5   Compute
      (21)  $\rho_4(\boldsymbol{\theta}', \boldsymbol{\theta}_{(t-1)}) = \min \left\{ 1, \frac{\pi^K(\boldsymbol{\theta}' | \mathbf{b}) r(\boldsymbol{\theta}_{(t-1)} | \boldsymbol{\theta}')}{\pi^K(\boldsymbol{\theta}_{(t-1)} | \mathbf{b}) r(\boldsymbol{\theta}' | \boldsymbol{\theta}_{(t-1)})} \right\}$ .
6   With probability  $\rho_4(\boldsymbol{\theta}', \boldsymbol{\theta}_{(t-1)})$ , set  $\boldsymbol{\theta}_{(t)} = \boldsymbol{\theta}'$  and  $\{\mathbf{x}_{(t,j)}\}_{j=1}^K = \{\mathbf{x}'_j\}_{j=1}^K$ , else set  $\boldsymbol{\theta}_{(t)} = \boldsymbol{\theta}_{(t-1)}$  and  $\{\mathbf{x}_{(t,j)}\}_{j=1}^K = \{\mathbf{x}_{(t-1,j)}\}_{j=1}^K$ .
7 end
```

All three algorithms require the computation of samples from the approximate conditional  $\hat{\pi}(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})$ , which is of the form

$$\mathbf{x} | \mathbf{b}, \boldsymbol{\theta} \sim \mathcal{N}(\hat{\mathbf{x}}_{\text{cond}}(\boldsymbol{\theta}), \hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta})),$$

where  $\hat{\mathbf{x}}_{\text{cond}}(\boldsymbol{\theta}) \equiv \mu \hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta}) \mathbf{A}^\top \mathbf{b}$ . In subsection 4.4, we tailor these results to a specific choice of  $\hat{\pi}(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})$ . The following quantity will be important in what follows:

$$(22) \quad w(\mathbf{x}, \boldsymbol{\theta}) \equiv \exp \left( -\frac{1}{2} \mathbf{x}^\top (\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta})) \mathbf{x} \right).$$

An expression for the moments of  $w(\mathbf{x}, \boldsymbol{\theta})$  can be computed analytically by using properties of Gaussian integrals and was established in [8]. For positive integers  $m$ ,

$$(23) \quad \mathbb{E}_{\hat{\pi}(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})}[w^m(\mathbf{x}, \boldsymbol{\theta})] = \frac{1}{M_m(\boldsymbol{\theta})},$$

where

$$(24) \quad M_m(\boldsymbol{\theta}) \equiv \frac{\exp \left( \frac{\mu^2}{2} \mathbf{b}^\top \mathbf{A} (\mathbf{M}_m^{-1}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta})) \mathbf{A}^\top \mathbf{b} \right)}{(\det \hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta}) \det \boldsymbol{\Gamma}_m(\boldsymbol{\theta}))^{1/2}}$$

with

$$\boldsymbol{\Gamma}_m(\boldsymbol{\theta}) = m(\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta})) + \hat{\boldsymbol{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta}).$$

Further results for  $M_m(\boldsymbol{\theta})$  can be derived if  $\hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta})$  is known explicitly. When  $\hat{\boldsymbol{\Gamma}}_{\text{cond}}(\boldsymbol{\theta})$  is constructed using the low-rank approach outlined in subsection 4.4 below, we have that  $M_m(\boldsymbol{\theta}) \geq 1$ .

**4.1. Analysis of the approximate one-block acceptance ratio.** Our first result derives a simplified version of the acceptance ratio of AOB (Algorithm 3) that is more amenable to interpretation.

**Proposition 1.** In Algorithm 3, let  $(\mathbf{x}, \boldsymbol{\theta})$  denote the current state of the AOB chain and let  $(\mathbf{x}', \boldsymbol{\theta}')$  be the proposed state. Then, the acceptance ratio simplifies to

$$\rho_2(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{z(\mathbf{x}', \boldsymbol{\theta}') \pi(\boldsymbol{\theta}' | \mathbf{b}) r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{z(\mathbf{x}, \boldsymbol{\theta}) \pi(\boldsymbol{\theta} | \mathbf{b}) r(\boldsymbol{\theta}' | \boldsymbol{\theta})} \right\},$$

where  $z(\mathbf{x}, \boldsymbol{\theta}) \equiv w(\mathbf{x}, \boldsymbol{\theta}) / M_1(\boldsymbol{\theta})$ .

**Proof.** It is clear that we only need to focus on the second term in the acceptance ratio, which we can rewrite as

$$\frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b}) \hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b}) r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b}) \hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b}) r(\boldsymbol{\theta}' | \boldsymbol{\theta})} = \frac{r(\boldsymbol{\theta} | \boldsymbol{\theta}') \pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b}) / \hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b})}{r(\boldsymbol{\theta}' | \boldsymbol{\theta}) \pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b}) / \hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})}.$$

Since the posterior distribution is the product of the conditional and the marginal, we have

$$(25) \quad \frac{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})}{\hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})} = \frac{\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b}) \pi(\boldsymbol{\theta} | \mathbf{b})}{\hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})}.$$

In the proof of [8, Proposition 1], it is shown that  $\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b}) / \hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b}) = z(\mathbf{x}, \boldsymbol{\theta})$ , so that we get

$$\frac{r(\boldsymbol{\theta} | \boldsymbol{\theta}') \pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b}) / \hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b})}{r(\boldsymbol{\theta}' | \boldsymbol{\theta}) \pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b}) / \hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})} = \frac{r(\boldsymbol{\theta} | \boldsymbol{\theta}') z(\mathbf{x}', \boldsymbol{\theta}') \pi(\boldsymbol{\theta}' | \mathbf{b})}{r(\boldsymbol{\theta}' | \boldsymbol{\theta}) z(\mathbf{x}, \boldsymbol{\theta}) \pi(\boldsymbol{\theta} | \mathbf{b})},$$

which gives the desired result. ■

The meaning of this result is clear if we combine (23) and [8, Proposition 2] to obtain

$$(26) \quad \mathbb{E}_{\hat{\pi}(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})} [z(\mathbf{x}, \boldsymbol{\theta})] = M_1^{-1}(\boldsymbol{\theta}) \mathbb{E}_{\hat{\pi}(\mathbf{x} | \mathbf{b}, \boldsymbol{\theta})} [w(\mathbf{x}, \boldsymbol{\theta})] = 1.$$

In other words, given the current state  $(\mathbf{x}, \boldsymbol{\theta})$  and a proposed state  $\boldsymbol{\theta}'$ , the ratio in Algorithm 3, when averaged over  $\hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b})$ , is approximately that of the one-block algorithm, provided  $z(\mathbf{x}, \boldsymbol{\theta}) \approx 1$ . Observe that if one takes  $\hat{\Gamma}_{\text{cond}}(\boldsymbol{\theta}) = \Gamma_{\text{cond}}(\boldsymbol{\theta})$ , so that  $z(\mathbf{x}, \boldsymbol{\theta}) = 1$ , then

$$\rho_2(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{\pi(\boldsymbol{\theta}' | \mathbf{b}) r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\pi(\boldsymbol{\theta} | \mathbf{b}) r(\boldsymbol{\theta}' | \boldsymbol{\theta})} \right\},$$

which is exactly the one-block acceptance ratio (15). Similarly, if  $\hat{\Gamma}_{\text{cond}}(\boldsymbol{\theta}) \approx \Gamma_{\text{cond}}(\boldsymbol{\theta})$  is a sufficiently accurate approximation, then  $z(\mathbf{x}, \boldsymbol{\theta}) \approx 1$ , and AOB will closely approximate the one-block algorithm.

**4.2. Analysis of the ABDA acceptance ratios.** Next, we discuss the acceptance ratio of ABDA (Algorithm 4). In the first stage, one-block is applied to  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$ , yielding the acceptance ratio (18) and target distribution  $\hat{\pi}(\boldsymbol{\theta} | \mathbf{b})$ . The analysis for the acceptance rate at the second stage is provided in Proposition 2.



**Proposition 2.** Let  $(\mathbf{x}, \boldsymbol{\theta})$  denote the current state of the ABDA chain and let  $(\mathbf{x}', \boldsymbol{\theta}')$  be the proposed state. With  $w(\mathbf{x}, \boldsymbol{\theta})$  defined in (22), the acceptance ratio at the second stage is

$$\rho_3(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{w(\mathbf{x}', \boldsymbol{\theta}')}{w(\mathbf{x}, \boldsymbol{\theta})} \right\}.$$

*Proof.* First, note that  $\hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta})$ , defined in (18), can be expressed as  $\min\{1, \gamma\}$ , where

$$\gamma \equiv \frac{\hat{\pi}(\boldsymbol{\theta}' | \mathbf{b})r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{\hat{\pi}(\boldsymbol{\theta} | \mathbf{b})r(\boldsymbol{\theta}' | \boldsymbol{\theta})}.$$

Similarly,  $\hat{\rho}_1(\boldsymbol{\theta}; \boldsymbol{\theta}') = \min\{1, \gamma^{-1}\}$ . But if  $\gamma > 0$ , then  $\min\{1, \gamma\} / \min\{1, \gamma^{-1}\} = \gamma$ . Therefore, the ratio  $\hat{\rho}_1(\boldsymbol{\theta}'; \boldsymbol{\theta}) / \hat{\rho}_1(\boldsymbol{\theta}; \boldsymbol{\theta}')$  simplifies to  $\gamma$ , from which it follows that

$$\frac{q(\mathbf{x}, \boldsymbol{\theta} | \mathbf{x}', \boldsymbol{\theta}')}{q(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{x}, \boldsymbol{\theta})} = \frac{\hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})\hat{\pi}(\boldsymbol{\theta}' | \mathbf{b})}{\hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b})\hat{\pi}(\boldsymbol{\theta} | \mathbf{b})},$$

where  $q$  is defined by (19). Thus, the ratio appearing in (17) simplifies to

$$\begin{aligned} \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})q(\mathbf{x}, \boldsymbol{\theta} | \mathbf{x}', \boldsymbol{\theta}')}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})q(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{x}, \boldsymbol{\theta})} &= \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})\hat{\pi}(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})\hat{\pi}(\boldsymbol{\theta} | \mathbf{b})}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})\hat{\pi}(\mathbf{x}' | \boldsymbol{\theta}', \mathbf{b})\hat{\pi}(\boldsymbol{\theta}' | \mathbf{b})} \\ &= \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})\hat{\pi}(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})}. \end{aligned}$$

Therefore, the acceptance ratio at the second stage is

$$(27) \quad \rho_3(\mathbf{x}', \boldsymbol{\theta}'; \mathbf{x}, \boldsymbol{\theta}) = \min \left\{ 1, \frac{\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})}{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})\hat{\pi}(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})} \right\}.$$

From this equation it is clear that we only need to focus on the ratio of the exact to approximate posterior densities  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b}) / \hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$ , which simplifies according to (25). Moreover, it is straightforward to verify that  $M_1(\boldsymbol{\theta}) = \hat{\pi}(\boldsymbol{\theta} | \mathbf{b}) / \pi(\boldsymbol{\theta} | \mathbf{b})$ , where  $M_1(\boldsymbol{\theta})$  is defined in (24). Substituting these two results into (27) and recalling that  $z(\mathbf{x}, \boldsymbol{\theta}) = w(\mathbf{x}, \boldsymbol{\theta}) / M_1(\boldsymbol{\theta})$ , we obtain the desired result. ■

In this analysis, (27) shows that Algorithm 4 amounts to generating a proposal from the approximate posterior distribution  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  in a Metropolis–Hastings independence sampler. Compared to Algorithm 3, we expect Algorithm 4 to have lower statistical efficiency but with much improved computational efficiency, since the posterior distribution  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  is only evaluated when a “good” sample has been drawn.

Proposition 2 states that the acceptance ratio at the second stage is close to 1 if  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  is a good approximation to  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$ . The value of the ABDA algorithm is seen by observing that regardless of the choice of proposal distribution for  $\boldsymbol{\theta}$ , the acceptance rate at the second stage will be high when  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  closely approximates the true target density  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$ . For instance, if a poor proposal density is used for  $\boldsymbol{\theta}$ , then many bad states  $\boldsymbol{\theta}'$  will likely be proposed, but they will be discarded without wasting the computational effort to evaluate  $\pi(\mathbf{x}', \boldsymbol{\theta}' | \mathbf{b})$ . Provided we construct  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  to closely approximate  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  for all  $(\mathbf{x}, \boldsymbol{\theta}) \in \text{supp } \pi$ , we can be confident that the true density will only be evaluated for states that have a high chance of being accepted. Further, using the approximate density in the second stage reduces the problem of tuning a Metropolis(–Hastings) algorithm to one of tuning the proposal  $r(\boldsymbol{\theta}' | \boldsymbol{\theta})$ , which is easier to do when  $\boldsymbol{\theta}$  is low-dimensional.

**4.3. Analysis of the pseudomarginal acceptance ratio.** We turn to the analysis of the pseudomarginal algorithm (Algorithm 5). It is worth recalling that for  $K = 1$ , the pseudomarginal algorithm reduces to AOB.

**Proposition 3.** *In Algorithm 5, the acceptance ratio simplifies to*

$$\rho_4(\boldsymbol{\theta}'; \boldsymbol{\theta}) = \min \left\{ 1, \frac{z_K(\boldsymbol{\theta}')\pi(\boldsymbol{\theta}' | \mathbf{b})r(\boldsymbol{\theta} | \boldsymbol{\theta}')}{z_K(\boldsymbol{\theta})\pi(\boldsymbol{\theta} | \mathbf{b})r(\boldsymbol{\theta}' | \boldsymbol{\theta})} \right\},$$

where  $z_K(\boldsymbol{\theta}) \equiv \frac{1}{K} \sum_{j=1}^K z(\mathbf{x}^j, \boldsymbol{\theta})$ .

*Proof.* The proof is similar to Proposition 1 and is omitted. ■

More insight can be obtained by considering the asymptotic behavior of  $z_K(\boldsymbol{\theta})$ . By (26) and the strong law of large numbers, it follows that, for fixed  $\boldsymbol{\theta}$ ,  $z_K(\boldsymbol{\theta}) \xrightarrow{\text{a.s.}} 1$  as  $K \rightarrow \infty$ . The interpretation is that the pseudomarginal algorithm exhibits similar behavior as AOB with respect to one-block. The difference is that, by taking  $K \rightarrow \infty$ , we attain almost sure convergence, as opposed to an average behavior. Further, by the central limit theorem and the fact that  $\mathbf{x}^1, \dots, \mathbf{x}^K$  are independent for fixed  $\boldsymbol{\theta}$ , we have that  $\sqrt{K}(z_K(\boldsymbol{\theta}) - 1) \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma_z^2(\boldsymbol{\theta}))$ , where

$$\sigma_z^2(\boldsymbol{\theta}) \equiv \text{Var}_{\hat{\pi}(\mathbf{x}|\boldsymbol{\theta}, \mathbf{b})}[z(\mathbf{x}, \boldsymbol{\theta})] = \frac{1}{M_1^2(\boldsymbol{\theta})} \left( \frac{1}{M_2(\boldsymbol{\theta})} - \frac{1}{M_1^2(\boldsymbol{\theta})} \right) < \infty \quad \forall \boldsymbol{\theta} \text{ (a.e.)}.$$

It follows from the central limit theorem that  $z_K(\boldsymbol{\theta})$  is  $\sqrt{K}$ -consistent for 1 [37], i.e.,

$$(28) \quad z_K(\boldsymbol{\theta}) - 1 = \mathcal{O}_p(K^{-1/2})$$

as  $K \rightarrow \infty$ . However, it is important to observe that  $\sigma_z(\boldsymbol{\theta})$  will be close to zero when the approximate distribution is close to the true distribution, in which case  $z_K(\boldsymbol{\theta})$  will be close to one with high probability even for small  $K$ .

This analysis shows the advantage of the pseudomarginal approach, namely, that we can still achieve desirable marginalization behavior in the presence of a poor approximation to the full conditional distribution of  $\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$ . If the approximation  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  is not very good, one can set  $K$  to be large to guarantee  $z_K(\boldsymbol{\theta}) \approx 1$ , at the expense of drawing repeated realizations from the approximating distribution. Otherwise, if the approximation  $\hat{\pi}(\mathbf{x}, \boldsymbol{\theta} | \mathbf{b})$  is good,  $z_K(\boldsymbol{\theta}) \approx 1$  even for small  $K$ . The trade-off is thus a large number of realizations from a poor approximation or few (one in the case of AOB) realizations from a quality approximation. Regardless, if  $z_K(\boldsymbol{\theta}) \approx 1$ , Algorithm 5 will closely approximate the one-block algorithm.

**4.4. Approximate posterior distribution using low-rank approximation.** We briefly review the low-rank approach for sampling from the conditional distribution  $\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{b})$  that was used in [8] and which previously appeared in [9, 10, 14, 15, 21, 48, 32]. First, recall that  $\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) = \mu \mathbf{A}^\top \mathbf{A} + \sigma \boldsymbol{\Gamma}_{\text{prior}}^{-1}$  and assume  $\boldsymbol{\Gamma}_{\text{prior}}^{-1} = \mathbf{L}^\top \mathbf{L}$ , so that

$$\boldsymbol{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) = \mathbf{L}^\top \left( \mu \mathbf{L}^{-\top} \mathbf{A}^\top \mathbf{A} \mathbf{L}^{-1} + \sigma \mathbf{I} \right) \mathbf{L}.$$

Next, define a rank- $k$  approximation as

$$(29) \quad \mathbf{L}^{-\top} \mathbf{A}^\top \mathbf{A} \mathbf{L}^{-1} \approx \mathbf{V}_k \boldsymbol{\Lambda}_k \mathbf{V}_k^\top,$$

where  $\mathbf{\Lambda}_k$  is diagonal with the  $k$  largest eigenvalues of  $\mathbf{L}^{-\top} \mathbf{A}^\top \mathbf{A} \mathbf{L}^{-1}$  and  $\mathbf{V}_k$  contains orthonormal columns. Combining the previous two equations, we obtain the approximate covariance

$$\hat{\mathbf{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta}) \equiv (\mu \mathbf{L}^\top \mathbf{V}_k \mathbf{\Lambda}_k \mathbf{V}_k^\top \mathbf{L} + \sigma \mathbf{\Gamma}_{\text{prior}}^{-1}) = \mathbf{L}^\top (\mu \mathbf{V}_k \mathbf{\Lambda}_k \mathbf{V}_k^\top + \sigma \mathbf{I}) \mathbf{L}.$$

Defining  $\hat{\mathbf{x}}_{\text{cond}}(\boldsymbol{\theta}) = \mu \hat{\mathbf{\Gamma}}_{\text{cond}}(\boldsymbol{\theta}) \mathbf{A}^\top \mathbf{b}$ , the approximate posterior distribution is obtained as

$$(30) \quad \hat{\pi}(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{b}) \propto f(\boldsymbol{\theta}) \exp \left( -\frac{\mu}{2} \mathbf{b}^\top \mathbf{b} + \frac{\mu^2}{2} \mathbf{b}^\top \mathbf{A} \hat{\mathbf{\Gamma}}_{\text{cond}}(\boldsymbol{\theta}) \mathbf{A}^\top \mathbf{b} - \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_{\text{cond}}(\boldsymbol{\theta})\|_{\hat{\mathbf{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta})}^2 \right),$$

where  $f(\boldsymbol{\theta}) = \mu^{M/2} \sigma^{N/2} \pi(\boldsymbol{\theta})$ . The marginal density is

$$(31) \quad \hat{\pi}(\boldsymbol{\theta} \mid \mathbf{b}) \propto \frac{f(\boldsymbol{\theta})}{\sqrt{\det(\hat{\mathbf{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta}))}} \exp \left( -\frac{\mu}{2} \mathbf{b}^\top \mathbf{b} + \frac{\mu^2}{2} \mathbf{b}^\top \mathbf{A} \hat{\mathbf{\Gamma}}_{\text{cond}}(\boldsymbol{\theta}) \mathbf{A}^\top \mathbf{b} \right).$$

To sample from the approximate conditional distribution  $\hat{\pi}(\boldsymbol{\theta} \mid \mathbf{b})$  we compute

$$(32) \quad \mathbf{x} = \hat{\mathbf{x}}_{\text{cond}}(\boldsymbol{\theta}) + \mathbf{G} \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

where  $\mathbf{G} := \sigma^{-1/2} \mathbf{L}^{-1} (\mathbf{I} - \mathbf{V}_k \hat{\mathbf{D}}_k \mathbf{V}_k^\top)$  and  $\hat{\mathbf{D}}_k = \mathbf{I} \pm (\mathbf{I} - \mathbf{D}_k)^{1/2}$ . It is shown in [8] that  $\hat{\mathbf{\Gamma}}_{\text{cond}} = \mathbf{G} \mathbf{G}^\top$ . Since  $\hat{\mathbf{D}}_k$  is diagonal and  $k \ll n$ , (32) provides a computationally cheap way of generating draws from  $\hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$ .

**Acceptance ratio analysis.** From [8, Proposition 2], the expression for  $w(\mathbf{x}, \boldsymbol{\theta})$  simplifies considerably. First,

$$(33) \quad \mathbf{\Gamma}_{\text{cond}}^{-1}(\boldsymbol{\theta}) - \hat{\mathbf{\Gamma}}_{\text{cond}}^{-1}(\boldsymbol{\theta}) = \mu \mathbf{L}^\top \left( \sum_{j=k+1}^N \lambda_j \mathbf{v}_j \mathbf{v}_j^\top \right) \mathbf{L},$$

implying  $w(\mathbf{x}, \boldsymbol{\theta}) = \exp(-\frac{1}{2} \sum_{j=k+1}^N \lambda_j (\mathbf{v}_j^\top \mathbf{L} \mathbf{x})^2)$ , which in turn implies [8, Theorem 1]

$$M_m(\boldsymbol{\theta}) = \exp \left( \frac{\mu^2}{2\sigma} \sum_{j=k+1}^n \frac{m\mu\lambda_j}{m\mu\lambda_j + \sigma} (\mathbf{b}^\top \mathbf{A} \mathbf{L}^{-1} \mathbf{v}_j)^2 \right) \prod_{j=k+1}^n \left( 1 + \frac{m\mu}{\sigma} \lambda_j \right)^{1/2}.$$

The key insight from this calculation is that  $z(\mathbf{x}, \boldsymbol{\theta}) = w(\mathbf{x}, \boldsymbol{\theta})/M_1(\boldsymbol{\theta})$  only contains eigenvalues that are discarded in the low-rank approximation. This means if the discarded eigenvalues are very small, AOB is very close to the one-block algorithm. Similarly, for the ABDA algorithm, this implies the acceptance ratio of the second stage is very high and for the pseudomarginal MCMC, the variance  $\sigma_z(\boldsymbol{\theta})$  is close to zero.

In practice, computing the truncated SVD can be computationally expensive. However, the low-rank decomposition can be approximately computed using a Krylov subspace approach [47] or using a randomized approach [29]. These approximate factorizations can be used instead to construct the approximating distribution, such as the conditional  $\hat{\pi}(\mathbf{x} \mid \mathbf{b}, \boldsymbol{\theta})$ , the marginal  $\hat{\pi}(\boldsymbol{\theta} \mid \mathbf{b})$ , and the full posterior distribution  $\hat{\pi}(\mathbf{x}, \mathbf{b} \mid \boldsymbol{\theta})$ . The details are given in [8].

Table 1

Summary of various computational costs in the AOB algorithm.

Component	Dominant cost	Additional cost
Pre-computation	$2k(T_{\mathbf{A}} + T_{\mathbf{L}^{-1}})$	$\mathcal{O}(nk^2)$
Sample	$T_{\mathbf{L}^{-1}}$	$\mathcal{O}(nk)$
Acceptance ratio	$T_{\mathbf{A}} + T_{\mathbf{L}}$	$\mathcal{O}(nk)$ .


**4.5. Computational cost.** We briefly review the computational cost of the three algorithms proposed in this paper. Denote the computational cost of forming the matrix-vector product (henceforth referred to as matvec) with  $\mathbf{A}$  by  $T_{\mathbf{A}}$ . Similarly denote the cost of a matvec with  $\mathbf{L}$  and  $\mathbf{L}^{-1}$  as  $T_{\mathbf{L}}$  and  $T_{\mathbf{L}^{-1}}$ , respectively. To simplify the analysis, we make two assumptions: first, the cost of the matvecs dominates the computational cost, and second, the cost of the transpose operations of the respective matrices is the same as that of the original matrix.

**AOB algorithm.** We first analyze the computational cost associated with AOB (Algorithm 3). There are three major sources of computational cost. In the offline stage, we precompute a low-rank approximation as in (29). By using a Krylov subspace solver or randomized SVD algorithm, the dominant cost is in computing the matvecs, i.e.,  $2k(T_{\mathbf{A}} + T_{\mathbf{L}^{-1}})$  floating point operations (flops), with an additional cost of  $\mathcal{O}(nk^2)$  flops. In the online stage, to generate a composite sample  $(\mathbf{x}', \boldsymbol{\theta}')$ , the major cost that depends on the dimension of the problem is due to (32) and can be quantified as  $T_{\mathbf{L}^{-1}}$  flops, with an additional  $\mathcal{O}(nk)$  cost. Finally, computing the acceptance ratio requires one evaluation of the full posterior distribution and one evaluation of the approximate posterior distribution. Evaluating the full posterior using subsection 2.1 requires  $T_{\mathbf{A}} + T_{\mathbf{L}}$  flops, with an additional cost  $\mathcal{O}(nk)$  flops, whereas evaluating the approximate posterior distribution using (30) only requires  $\mathcal{O}(nk)$  flops. These computational costs are summarized in Table 1.

**ABDA algorithm.** The cost of ABDA, Algorithm 4, is the same as that of AOB, Algorithm 3, with only one notable exception. The extra step of evaluating the approximate marginal distribution (31) requires an additional  $2T_{\mathbf{L}^{-1}} + \mathcal{O}(nk)$  flops. However, while the per-iteration cost of ABDA is comparable to that of AOB, the overall cost of ABDA can be lower. The reason is that the acceptance ratio involving the full posterior distribution is only evaluated in the second stage, unlike AOB in which the posterior distribution is evaluated at every iteration. The computational speedup has been demonstrated in numerical experiments (see section 5). Estimates of the computational speedup can be obtained by following the approach in [16], but we omit this discussion.

**Pseudomarginal algorithm.** The cost of the pseudomarginal approach (Algorithm 5) is the same as Algorithm 3 with two exceptions:  $K$  samples need to be generated which cost  $K[T_{\mathbf{L}^{-1}} + \mathcal{O}(nk)]$  flops, and to evaluate the acceptance ratio, we need  $K[T_{\mathbf{A}} + T_{\mathbf{L}} + \mathcal{O}(nk)]$  flops. The benefits of the pseudomarginal algorithm, i.e., more effective marginalization, have to be weighed against the additional cost of generating the samples and evaluating the acceptance ratio.

**5. Numerical experiments.** Throughout this section, the proposal distribution for  $\boldsymbol{\theta}$  is taken to be the adaptive Metropolis proposal [28] using a lognormal proposal for  $\boldsymbol{\theta}$ . For the prior and noise precision parameters, we assign a Gamma prior,  $\text{Gamma}(1, 10^{-4})$ , based on



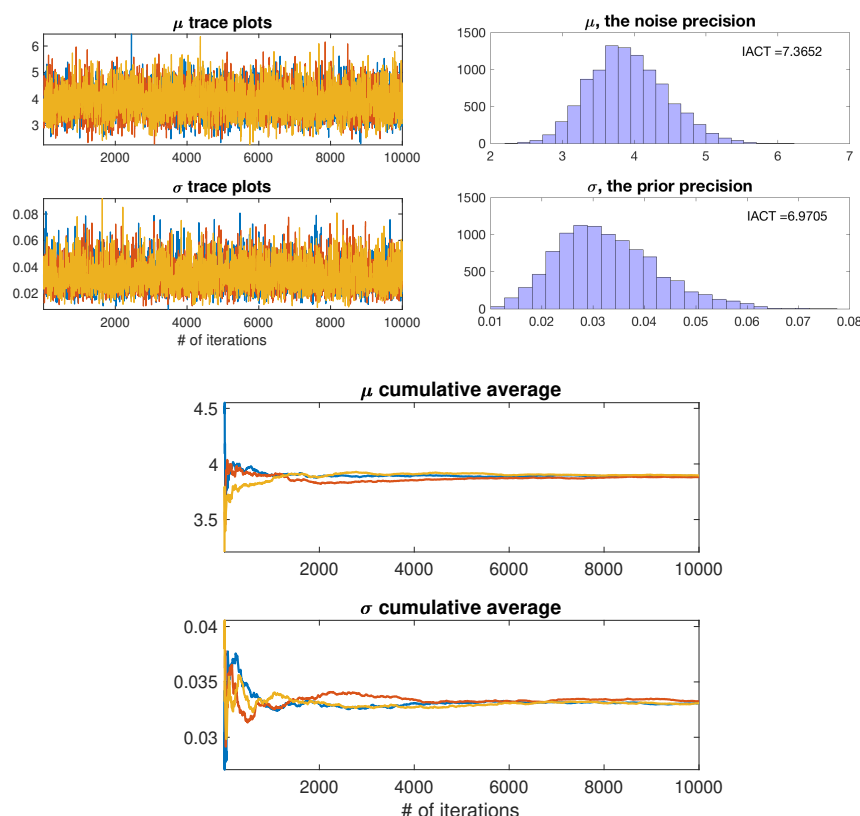
**Figure 2.** (left) The true solution  $\mathbf{x}_{\text{true}}$  and the blurred vector  $\mathbf{b}$  with 2% Gaussian noise to simulate measurement error. (center) The posterior mean superimposed with the pointwise 95% credible bounds. (right) The spectrum of the matrix  $\mathbf{L}^{-\top} \mathbf{A}^{\top} \mathbf{A} \mathbf{L}^{-1}$ .

the recommendation of [5]. If a priori knowledge is available regarding the distribution of  $\boldsymbol{\theta}$ , this can be incorporated into the hyperpriors. It is also worth mentioning that besides the Gamma priors, other choices for the hyper-priors can also be made, e.g., proper Jeffreys. See [8] for a detailed discussion. Besides these parameters, no other tuning parameters were necessary.

**5.1. One-dimensional deblurring.** We take a simple one-dimension deblurring example to illustrate our algorithms. This application arises from the discretization of a Fredholm integral equation of the first kind as in (11). The details of this application are given elsewhere and we refer the reader to them [3]. For the problem size we choose  $N = 128$  and add Gaussian noise with variance  $0.01^2 \|\mathbf{A} \mathbf{x}_{\text{true}}\|_2^2$  to simulate measurement error, where  $\mathbf{x}_{\text{true}}$  is the true vector which is assumed to be known. We model smoothness on  $\mathbf{x}$  a priori by taking  $\Gamma_{\text{prior}}^{-1} = -\Delta$ , where  $-\Delta$  is the discrete Laplacian with Dirichlet boundary conditions [35].

A low-rank approximation with  $k = 35$  computed using the “exact” SVD was used to define the approximate posterior distribution. We run 20,000 iterations of the AOB algorithm. The first 10,000 samples were considered to be the burn-in period and were discarded. The left panel of Figure 2 shows the true vector  $\mathbf{x}_{\text{true}}$  and the blurred vector  $\mathbf{b}$  with the simulated measurement noise. The right panel of the same figure shows the approximate posterior mean superimposed on the true vector  $\mathbf{x}_{\text{true}}$ ; also plotted are the curves corresponding to the pointwise 95% credible bounds. The plots show that the true image is mostly contained within the credible intervals, thereby providing a measure of confidence in the reconstruction process.

**Diagnostics for convergence.** We now provide some limited diagnostics to assess the convergence of the algorithm. In the upper left panel of Figure 3, we display the trace plots of the precision parameters  $\mu$  and  $\sigma$  for three different chains with different initializations. It is seen that the chains appear to converge. Furthermore, the multivariate potential scale reduction factor [7] for the  $\mathbf{x}$  chain is 1.06, which is within the recommend rule-of-thumb range (less than 1.1). The top right panel of the same figure shows the histograms of the precision parameters. The (two-sided) Geweke test [4] applied to the resulting  $\mu$ - and  $\sigma$ -chains yield  $p$ -values both greater than 0.98, meaning that there is no reason to believe the chains are out of equilibrium.




**Figure 3.** The (top left) trace plots and (top right) histograms of the  $\mu$  and  $\sigma$  chains for the one-dimensional deblurring example using the AOB algorithm. (bottom) Cumulative averages of the  $\mu$  and  $\sigma$  chains against the number of iterations. The target rank is  $k = 50$ .

Further support for convergence is obtained by monitoring the cumulative averages of the  $\mu$  and  $\sigma$ -chains plotted in the bottom panel of Figure 3, from which it is readily seen that the cumulative averages converge.

In the right panel of Figure 1 (in section 2), we show the empirical autocorrelation function of one of the  $\sigma$  chains. As mentioned earlier, the autocorrelation decays rapidly. While we do not plot the autocorrelation of the  $\mu$  chain, we observed that it had similar behavior. The integrated autocorrelation time (IACT) was 6.97 for the  $\sigma$ -chain. The acceptance rate was 33.8% and the effective sample size (ESS) was 1434.61. Here, the ESS [36] is defined as  $\text{ESS} = N_s / \tau_\sigma$ , where  $N_s$  is the Monte Carlo sample size and  $\tau_\sigma$  denotes IACT of the Markov chain.

**Comparing different methods.** We now compare all three algorithms proposed in section 3. The setup of the problem is the same as the previous experiment. The previous analysis shows that the AOB algorithm does a good job decorrelating the  $\sigma$ -chain from the  $\mathbf{x}$  chain. We repeat this experiment for the other two methods as well. Figure 4 plots the autocorrelation functions of the  $\sigma$ -chains produced by the other two methods, ABDA and PM with  $K = 5$ . As is readily seen, both methods successfully reduce the autocorrelation in the  $\sigma$ -chains. Although we do not plot the  $\mu$ -chains, they behave similarly to the  $\sigma$ -chains.





**Figure 4.** Plots of the empirical autocorrelation functions for the  $\sigma$  chains in the one-dimensional deblurring example: (left) ABDA and (right) PM with  $K = 5$ .

**Table 2**

IACT and ESS for the  $\sigma$  chains produced by the proposed algorithms for the one-dimensional deblurring example.

	IACT $\sigma$	ESS
AOB	6.97	1434.61
ABDA	7.14	1400.90
PM $K = 5$	6.90	1448.38

In Table 2, we compare the various methods proposed here. The ESS of the PM algorithm is marginally better than ABDA or AOB. This is because PM has a lower IACT. This is not surprising since the PM uses additional samples to marginalize the posterior distribution. More details regarding the PM method are given in the next experiment. Between the AOB and ABDA algorithms, AOB has a higher ESS and is more statistically efficient. In principle, though, ABDA is much less computationally expensive since the forward operator  $\mathbf{A}$  is evaluated less frequently. The difference is more pronounced in the inverse heat equation example in which the forward model is more expensive to evaluate.

**Effect of target rank.** Here we investigate the effect of the target rank on the acceptance rate and the statistical efficiencies of the samplers. The model is identical to that in the first experiment with the exception that we change the target rank  $k$  that controls the accuracy of the approximate posterior distribution. We vary the target rank  $k$  from 20 to 40 in increments of 5. The results are displayed in Table 3. When the target rank  $k$  is below 25, we find that the acceptance rate is very close to zero, indicating that the low-rank approximation is not sufficiently accurate. On the other hand, when the target rank  $k$  is 25 and above, the acceptance rate is close to 33%. It is also seen that increasing the target rank does not substantially increase the acceptance rate. This suggests that we have successfully approximated the posterior distribution to the point that the only limitation is the proposal distribution for  $\theta$ ; effectively, the dimensionality has been reduced from 128 to 25 which is roughly a factor of 5. Similar results are observed for the ABDA algorithm as well, shown in the right panel of Table 3.

**Table 3**

Summary of the effect of increasing rank for the one-dimensional deblurring example. “A.R.” refers to acceptance rate; the number as a suffix refers to the stage. In general, an increase in the target rank increases the acceptance ratio and the ESS.

Rank	ESS	A.R.	Rank	ESS	A.R. 1	A.R. 2
20	259.38	0.1123	20	7.67	0.0048	0.1465
25	1353.71	0.3318	25	1056.31	0.4988	0.7017
30	1480.83	0.3383	30	1326.67	0.5024	0.7242
35	1357.73	0.3396	35	1139.74	0.4997	0.7264
40	1357.73	0.3396	40	1139.74	0.4997	0.7264

(a) AOB

(b) ABDA

**Table 4**

Effect of the importance sample size  $K$  on the  $\sigma$ -chains produced by the pseudomarginal approach in the one-dimensional deblurring example. The target rank is fixed to be  $k = 20$ .

$K$	ESS	IAC $\sigma$	CPU time [s]	CES
1	282.49	32.00	43.62	0.1544
5	657.52	12.29	54.56	0.0830
10	1019.37	9.81	65.87	0.0646
50	1329.79	7.52	154.52	0.1162

**Effect of increased importance sample size.** In this experiment, we explore the effect of the importance sample size  $K$  on the statistical efficiency in the PM algorithm. In addition to the ESS and the IAC of the  $\sigma$ -chain, we report the CPU time in seconds, and the computational cost per effective sample (CES), which is the ratio of the CPU time to the ESS. [Proposition 3](#) shows that there are two ways to make the PM algorithm closer to one-block: by increasing the target rank that defines the approximate distribution, or by increasing the number of samples  $K$ . We take the target rank to be  $k = 20$ . From the previous experiment it is seen that the samples have high autocorrelation and the resulting ESS is poor. [Table 4](#) shows that by increasing the number of samples  $K$ , the PM approach seems to be performing better as evident in the increase in the ESS. However, the CPU time also increases considerably since each step in the algorithm is more expensive, with increasing  $K$ . Indeed, when  $K$  increases by a factor of 10, the ESS merely doubles. On the other hand, from [Table 3](#) (corresponding to PM with  $K = 1$ ), increasing the target rank  $k$  by 10 substantially increases the ESS, but is far less expensive. However, when the spectrum of the prior-preconditioned Hessian is flat, increasing the target rank may not improve the ESS. On the other hand, [Proposition 3](#) shows that increasing the sample size  $K$  will have the desired effect.

**5.2. A PDE-based example.** In this application, we consider the inversion of the two-dimensional initial state in the heat equation. Given the initial state  $u_0$ , we solve

$$\begin{aligned} u_t - \kappa \Delta u &= 0 && \text{in } \mathcal{D} \times [0, T], \\ u(x, t) &= 0 && \text{on } \partial \mathcal{D} \times [0, T], \\ u(x, 0) &= u_0 && \text{in } \mathcal{D}. \end{aligned}$$

The domain is taken to be  $\mathcal{D} = [0, 2] \times [0, 1]$ . In the present experiment the diffusion coefficient




Figure 5. True initial-state (left) and observation points (right) for the inverse heat equation example.

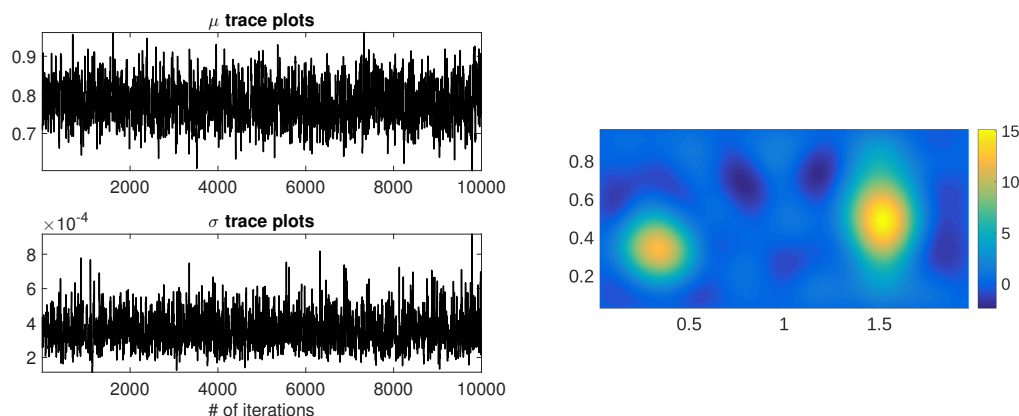


Figure 6. AOB results for the inverse heat equation example. Left:  $\mu$  and  $\sigma$  trace plots; right: posterior mean estimate of the initial state. The Monte Carlo sample size was  $10^4$ .

was set to  $\kappa = .001$  and we used a final simulation time of  $T = 5$ . The inverse problem seeks to use point measurements of  $u(\cdot, T)$  to reconstruct the initial state. The assumed true initial state used to simulate the data is shown in Figure 5 (left). We suppose that measurements are taken at an array of 512 observation points depicted in Figure 5 (right).

We used a Gaussian prior with covariance operator  $\mathcal{A}^{-2}$ , where  $\mathcal{A}$  is the Laplacian with zero Dirichlet boundary condition. The data were generated by adding 10% Gaussian noise to the measurements. The problem is discretized with a  $64 \times 32$  grid. In the inverse problems, we seek to reconstruct a discretized initial state in  $\mathbb{R}^n$  with  $n = 63 \times 31 = 1953$ .

Let  $\mathbf{A}$  denote the discretized forward operator that is obtained by composing the PDE solution operator and the observation operator that extracts solution values at the measurement points. Owing to the fast decay of singular values of  $\mathbf{A}$ , we can use a low-rank approximation. In the present example, a rank-70 approximation was found to provide sufficient accuracy.

**Performance of AOB and ABDA.** As a first experiment, we apply the AOB method to the present inverse problem. Our MCMC implementation is run for  $2 \times 10^4$  iterations; we retain  $10^4$  samples and discard the rest as part of the burn-in period. In Figure 6, we provide the trace plots of the  $\mu$  and  $\sigma$  chains (left), along with approximate posterior mean estimate of the initial state (right). We also report the empirical autocorrelation functions for the  $\mu$  and  $\sigma$  chains in Figure 7. The acceptance rate for the algorithm was approximately 30%.

Next, we compare the performance of AOB with that of ABDA. The results of our numerical experiments are summarized in Figure 8. We note that both methods produce essentially




Figure 7. The empirical autocorrelation function computed using AOB in the inverse heat equation example.

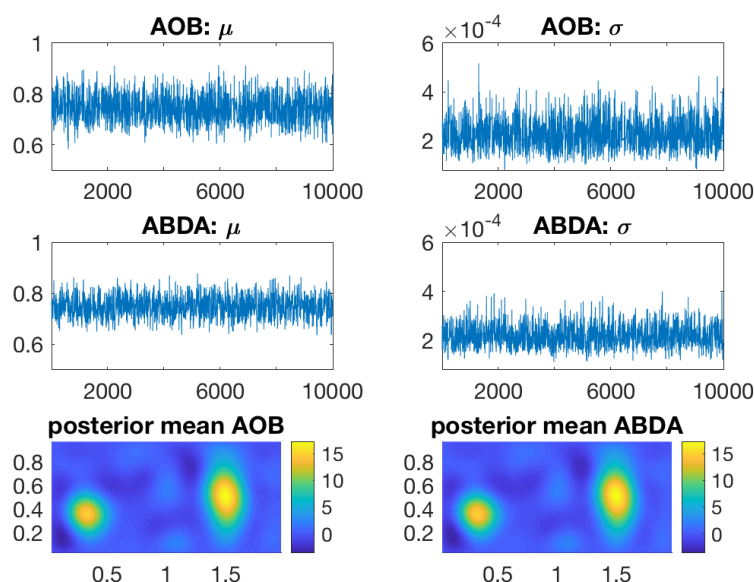


Figure 8. Statistical inversion results for the heat equation example, using AOB and ABDA methods.

the same results. However, Table 5 illustrates the reduction in computation time facilitated by ABDA. This application suggests that ABDA may be preferable to AOB when the forward problem is computationally expensive.

**6. Conclusions.** This paper focuses on the problem of sampling from a posterior distribution that arises in hierarchical Bayesian inverse problems. We restrict ourselves to linear inverse problems with Gaussian measurement error, and we assume a Gaussian prior on the unknown quantity of interest. The hierarchical model arises when the precision parameters for both the measurement error and prior are assigned prior distributions. Gibbs sampling

Table 5

Sampling metrics for AOB and ABDA methods for the inverse heat equation example.

Metric	AOB	ABDA
Time (sec)	819.89	622.19
ESS	1096.46	933.37
IACT	9.12	10.71

[3] and low-rank independence sampling [8] have been proposed for this type of problem, but the statistical efficiency of the resulting MCMC chains produced by either approach decreases as the dimension of the state increases. We discuss this phenomenon in detail and consider a solution using marginalization-based methods [22, 33, 45]. Marginalization can be just as computationally prohibitive for sufficiently large-scale problems. Thus, we combine the low-rank techniques with the marginalization-based approaches to propose three MCMC algorithms that are computationally feasible for larger-scaled problems. The first of these new MCMC methods is a direct extension of the one-block algorithm [45]; the second is an extension of the delayed acceptance algorithm [13]; and the last is an extension of the pseudomarginal algorithm [2]. We test and compare the performances of these three methods on two test cases, demonstrating that they all work reasonably well. We also offer suggestions on when one algorithm might be preferable over another based on, e.g., the spectrum of the prior-preconditioned Hessian matrix. Future work will consider extensions to nonlinear forward models and non-Gaussian priors on the unknown quantity of interest.

## REFERENCES

- [1] S. AGAPIOU, J. M. BARDSLEY, O. PAPASPILIOPOULOS, AND A. M. STUART, *Analysis of the Gibbs sampler for hierarchical inverse problems*, SIAM/ASA J. Uncertain. Quantif., 2 (2014), pp. 511–544.
- [2] C. ANDRIEU AND G. O. ROBERTS, *The pseudo-marginal approach for efficient Monte Carlo computations*, Ann. Statist., 37 (2009), pp. 697–725.
- [3] J. M. BARDSLEY, *MCMC-based image reconstruction with uncertainty quantification*, SIAM J. Sci. Comput., 34 (2012), pp. A1316–A1332.
- [4] J. M. BARDSLEY, *Computational Uncertainty Quantification for Inverse Problems*, SIAM, Philadelphia, 2018.
- [5] J. M. BARDSLEY, M. HOWARD, AND J. G. NAGY, *Efficient MCMC-based image deblurring with Neumann boundary conditions*, Electron. Trans. Numer. Anal., 40 (2013), pp. 476–488.
- [6] S. BROOKS, A. GELMAN, G. JONES, AND X.-L. MENG, *Handbook of Markov Chain Monte Carlo*, CRC Press, Boca, Raton, FL, 2011.
- [7] S. P. BROOKS AND A. GELMAN, *General methods for monitoring convergence of iterative simulations*, J. Comput. Graph. Statist., 7 (1998), pp. 434–455.
- [8] D. A. BROWN, A. K. SAIBABA, AND S. VALLÉLIAN, *Low-rank independence samplers in hierarchical Bayesian inverse problems*, SIAM/ASA J. Uncertain. Quantif., 6 (2018), pp. 1076–1100.
- [9] T. BUI-THANH, O. GHATTAS, J. MARTIN, AND G. STADLER, *A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion*, SIAM J. Sci. Comput., 35 (2013), pp. A2494–A2523.
- [10] D. CALVETTI AND E. SOMERSALO, *Priorconditioners for linear systems*, Inverse Problems, 21 (2005), 1397.
- [11] D. CALVETTI AND E. SOMERSALO, *Introduction to Bayesian Scientific Computing*, Springer, New York, 2007.
- [12] D. CALVETTI AND E. SOMERSALO, *Hypermmodels in the Bayesian imaging framework*, Inverse Problems, 24 (2008), 034013.

- [13] J. A. CHRISTEN AND C. FOX, *Markov chain Monte Carlo using an approximation*, J. Comput. Graph. Statist., 14 (2005), pp. 795–810.
- [14] J. CHUNG AND M. CHUNG, *Computing optimal low-rank matrix approximations for image processing*, in Proceedings of the Asilomar Conference on Signals, Systems and Computers, 2013, 670674.
- [15] J. CHUNG AND M. CHUNG, *An efficient approach for computing optimal low-rank regularized inverse matrices*, Inverse Problems, 30 (2014), 114009.
- [16] T. CUI, C. FOX, AND M. J. O’SULLIVAN, *Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm*, Water Resour. Res., 47 (2011).
- [17] T. CUI, K. LAW, AND Y. M. MARZOUK, *Dimension-independent likelihood-informed MCMC*, J. Comput. Phys., 304 (2016), pp. 109–137.
- [18] G. DA PRATO, *An Introduction to Infinite-Dimensional Analysis*, Universitext, Springer, New York, 2006.
- [19] G. DA PRATO AND J. ZABCZYK, *Second Order Partial Differential Equations in Hilbert Spaces*, London Math. Soc. Lecture Note Ser. 293, Cambridge University Press, Cambridge, 2002.
- [20] M. DASHTI AND A. M. STUART, *The Bayesian approach to inverse problems*, in Handbook of Uncertainty Quantification, R. Ghanem, D. Higdon, and H. Owhadi, eds., Springer, New York, 2017, pp. 311–428.
- [21] H. FLATH, L. WILCOX, V. AKCELIK, J. HILL, B. VAN BLOEMEN WAANDERS, AND O. GHATTAS, *Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations*, SIAM J. Sci. Comput., 33 (2011), 407432.
- [22] C. FOX AND R. A. NORTON, *Fast sampling in a linear-Gaussian inverse problem*, SIAM/ASA J. Uncertain. Quantif., 4 (2016), pp. 1191–1218.
- [23] A. E. GELFAND AND A. F. M. SMITH, *Sampling-based approaches to calculating marginal densities*, J. Amer. Statist. Assoc., 85 (1990), pp. 398–409.
- [24] A. GELMAN, *Prior distributions for variance parameters in hierarchical models*, Bayesian Anal., 1 (2006), pp. 515–533.
- [25] A. GELMAN, J. B. CARLIN, H. S. STERN, D. B. DUNSON, A. VEHTARI, AND D. B. RUBIN, *Bayesian Data Analysis*, 3rd ed., Texts Statist. Sci. Ser., CRC Press, Boca Raton, FL, 2014.
- [26] S. GEMAN AND D. GEMAN, *Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images*, IEEE Trans. Pattern Anal. Mach. Intell., 6 (1984), pp. 721–741.
- [27] C. GILAVERT, S. MOUSSAOUI, AND J. IDIER, *Efficient Gaussian sampling for solving large-scale inverse problems using MCMC*, IEEE Trans. Signal Process., 63 (2015), pp. 70–80.
- [28] H. HAARIO, E. SAKSMAN, AND J. TAMMINEN, *An adaptive Metropolis algorithm*, Bernoulli, 7 (2001), pp. 223–242.
- [29] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [30] W. K. HASTINGS, *Monte Carlo sampling methods using Markov chains and their applications*, Biometrika, 57 (1970), pp. 97–109.
- [31] D. HIGDON, *A primer on space-time modelling from a Bayesian perspective*, in Statistical Methods for Spatio-Temporal Systems, Finkenstadt, Held, and Isham, eds., CRC Press, Boca Raton, FL, 2006, pp. 218–277.
- [32] T. ISAAC, N. PETRA, G. STADLER, AND O. GHATTAS, *Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet*, J. Comput. Phys., 296 (2015), pp. 348–368.
- [33] K. JOYCE, J. BARDSLEY, AND A. LUTTMAN, *Point spread function estimation in X-ray imaging with partially collapsed Gibbs sampling*, SIAM J. Sci. Comput., 40 (2018), pp. B766–B787.
- [34] J. KAIPIO, V. KOLEHMAINEN, E. SOMERSALO, AND M. VAUHKONEN, *Statistical inversion and Monte Carlo sampling methods in electrical impedance tomography*, Inverse Problems, 16 (2000), 14871522.
- [35] J. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*. Springer, New York, 2005.
- [36] R. E. KASS, B. P. CARLIN, A. GELMAN, AND R. NEAL, *Markov chain Monte Carlo in practice: A roundtable discussion*, Amer. Statist., 52 (1998), pp. 93–100.
- [37] E. LEHMANN AND G. CASELLA, *Theory of Point Estimation*, 2nd ed., Springer, New York, 1998.
- [38] J. S. LIU, *Monte Carlo Strategies in Scientific Computing*, Springer, New York, 2004.
- [39] J. MARTIN, L. C. WILCOX, C. BURSTEDDE, AND O. GHATTAS, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM J. Sci. Comput., 34 (2012), pp. A1460–A1487.



- [40] N. METROPOLIS, A. W. ROSENBLUTH, M. N. ROSENBLUTH, A. H. TELLER, AND E. TELLER, *Equation of state calculations by fast computing machines*, J. Chem. Phys., 21 (1953), pp. 1087–1092.
- [41] G. NICHOLLS AND C. FOX, *Prior modelling and posterior sampling in impedance imaging*, in Bayesian Inference for Inverse Problems, Proc. SPIE 3459, 1998, pp. 116–127.
- [42] M. PARNO AND Y. M. MARZOUK, *Transport map accelerated Markov chain Monte Carlo*, SIAM/ASA J. Uncertain. Quantif., 6 (2018), pp. 645–682.
- [43] N. PETRA, J. MARTIN, G. STADLER, AND O. GHATTAS, *A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet inverse problems*, SIAM J. Sci. Comput., 36 (2014), pp. A1525–A1555.
- [44] C. ROBERT AND G. CASELLA, *Monte Carlo Statistical Methods*, 2nd ed., Springer, New York, 2004.
- [45] H. RUE AND L. HELD, *Gaussian Markov Random Fields: Theory and Applications*, Monogr. Statist. Appl. Probab. 104, Chapman & Hall/CRC, Boca Raton, FL, 2005.
- [46] J. G. SCOTT AND J. O. BERGER, *An exploration of aspects of Bayesian multiple testing*, J. Statist. Plann. Inference, 136 (2006), pp. 2144–2162.
- [47] H. D. SIMON AND H. ZHA, *Low-rank matrix approximation using the Lanczos bidiagonalization process with applications*, SIAM J. Sci. Comput., 21 (2000), pp. 2257–2274.
- [48] A. SPANTINI, A. SOLONEN, T. CUI, J. MARTIN, L. TENORIO, AND Y. MARZOUK, *Optimal low-rank approximations of Bayesian linear inverse problems*, SIAM J. Sci. Comput., 37 (2015), pp. A2451–A2487.
- [49] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numer., 19 (2010), pp. 451–559.
- [50] L. TIERNEY, *Markov chains for exploring posterior distributions*, Ann. Statist., 22 (1994), pp. 1701–1728.